Astronomy & Astrophysics manuscript no. gaiaplanets
April 11, 2017
©ESO 2017

Astrometry and exoplanets in the Gaia era: a Bayesian approach to detection and parameter recovery

P. Ranalli[1,2], D. Hobbs[1], and L. Lindegren[1]

[1] Lund Observatory, Box 43, 22100 Lund, Sweden
e-mail: piero@astro.lu.se
[2] INAF – Osservatorio Astronomico di Bologna, via Ranzani 1, 40127 Bologna, Italy

April 11, 2017

arXiv:1704.02493v1 [astro-ph.EP] 8 Apr 2017

ABSTRACT

The Gaia mission is expected to make a significant contribution to the knowledge of exoplanet systems, both in terms of their number and of their physical properties. We develop Bayesian methods and detection criteria for orbital fitting, and revise the detectability of exoplanets in light of the in-flight properties of Gaia. We start by determining what kind of exoplanets can be detected by Gaia, as a function of signal/noise ratio (SN), period, and distance from the Solar system. Then, we generate four sets of simulations of Gaia data for exoplanet systems over a grid of SN, orbital period, and eccentricity; one couple of sets for the nominal mission length (5 yr), and one couple for an extended mission length (10 yr); one couple of sets contains no planets, and one couple contains one planet. We define two fitting models: an astrometric model containing no planet, and a Keplerian model containing one planet. We fit our simulations using `Stan`, an off-the-shelf software for Bayesian inference. We investigate the detection rate according to four different indicators: the Akaike information criterion (AIC), the Bayesian information criterion (BIC), the widely-applicable information criterion (WAIC), and the $\Delta\chi^2$. We calibrate detection thresholds that select essentially no false positives, and can detect 70–75% of planets at high SN when convergence of the Markov chains is attained. We find that all four detection criteria can detect planets with SN $\gtrsim$ 1–1.5 for the 5 yr mission, and 0.7–1 for the 10 yr mission. This sets the maximum distance to which a planet is detectable to $\sim$ 160 pc and $\sim$ 8 pc for a Jovian and Neptunian planet, respectively, assuming an extended 10 yr mission, a 4 au semi-major axis, and a 0.5 $M_\odot$ star. We show how the orbital parameters are recovered for both mission lengths, and present a summary of the characteristic size of the error bars of each parameter. The period is the orbital parameter that can be determined with the best accuracy, with a median relative error of the order of $\pm$1.3%. The angular dimension of the semi-major axis of the orbit can be recovered with a median relative error of $\pm$17%. The eccentricity can also be recovered with a median absolute error of $\pm$0.14. Moreover, considering systems for which the Markov chains did not converge, we find that the WAIC can operate a much cleaner selection of systems with a good fit than the AIC or the BIC. We conclude that the Bayesian framework offers effective tools for modelling astrometric data of stars with exoplanets, both in terms of available software, and of theoretical methods. Model selection based on information criteria is an effective way to decide whether a planet is detected; with the WAIC being better suited then the AIC or the BIC in some circumstances.

Key words. methods: statistical – methods: numerical – astrometry – celestial mechanics – techniques: miscellaneous – (stars:) planetary systems

1. Introduction

The detection of exoplanets with astrometric techniques is simple in principle. Although an exoplanet is usually too faint to be observed with current instruments, one could infer its existence from the motion of its host star around the centre of mass of the star-planet system. In the past, several claims of detections have been made and later disproven starting from Jacob (1855); see Quirrenbach (2010) for a brief historical account. The difficulty lies in the small magnitude of the host star astrometric shift, which even in the most favourable cases is $\lesssim$ 1 mas (see e.g. Sect. 3).

The astrometric satellite Hipparcos (ESA 1997) could measure the along-scan position of a star with an error of a few mas. Several studies tried to derive orbital constraints for candidate planets suggested by radial velocity searches (Mazeh et al. 1999; Zucker & Mazeh 2000; Gatewood et al. 2001; Han et al. 2001) but fitting orbital models to Hipparcos data hardly improved the astrometric fit and led instead to spurious features (Pourbaix 2001).

The first sensible contribution by astrometry to the knowledge of an exoplanet system orbits happened with the fine guidance sensors (FGS) instrument on-board the Hubble space telescope (e.g., McArthur et al. 2010). However, the FGS is not designed for surveys so it is unlikely that large numbers of new planets could be discovered with it.

The Gaia mission (Gaia Collaboration 2016) offers a $\sim$ 30-fold improvement in astrometric precision with respect to Hipparcos, and can therefore finally allow the astrometric detection of planets to become routine. The Gaia capabilities for planet detection were explored by Casertano et al. (2008, hereafter C08) and later revised by Perryman et al. (2014, hereafter P14). The number of detectable planets was estimated by P14 to be $\sim 21,000$ for a 5 yr mission,

after considering a model of the Galaxy population and fitting planetary orbits with the least-squares method.

In this paper, we consider the application of Bayesian methods to the problem of planetary orbit fitting. Among the reasons why Bayesian methods should be considered, we focus on two: the ability of obtaining joint errors for all parameters of interest, and the availability of detection metrics which overcome the difficulties that arise for calculating goodness-of-fit for non-linear models. Bayesian methods are already in use for the Gaia data processing, and thus we expect our results to better anticipate the outcome of future Gaia data releases. We also consider $\Delta\chi^2$ as a detection metric: we show the need for empirically calibrating its significance scale and compare it to the other detection metrics.

Similarly to C08 and at variance with P14, we focus more on the detectability as a function of the orbital characteristics, than on the number of detectable planets in the Solar neighbourhood. While C08 was published five years before the launch of Gaia and therefore it reflects pre-launch expectations, in this paper we use the current knowledge of the satellite to reassess what fraction of planets can be detected and how well their orbital parameters can be recovered.

Because in this paper we focus on developing methods, we only consider single-planet systems. This is a reasonable approximation for real-world planetary systems where only one planet has a signal/noise ratio large enough. Double-planet and more complicated systems will be considered in a follow-up paper.

The structure of this paper is as follows. In Sect. 2 we describe our adopted formalism for the description of Keplerian orbits in an astrometry setting. In Sect. 3 we consider the expected signal/noise ratio for a one-planet system. In Sect. 4 we describe our simulations. In Sect. 5 we consider Bayesian methods and detection metrics. In Sect. 6 we present our results for the detection rate, and for the recovery of each parameter. In Sect. 7 we discuss the average error with which a parameter can be recovered, and the issue of non-converging Markov chains. Finally, in Sect. 8, we present our conclusions.

2. Astrometry and Keplerian orbits

2.1. Formalism

An exoplanet orbit can be determined in the same way as that of astrometric binary stars, where only one element of the binary system is visible. Here we summarise the method, following Binnendijk (1960), Quirrenbach (2010), and Perryman (2011), to which we refer for any further detail.

Let us consider a system made of a star and one planet. In an inertial reference frame, both bodies follow elliptical orbits around the system's centre of mass. The semimajor axis of the star's orbit is

$$a_* = a_p M_p / M_* \tag{1}$$

where a_p is the semimajor axis of the planet orbit, M_p is the planet mass, and M_* is the stellar mass. The astrometric signature[1] υ is the angular size of the stellar orbit

$$\upsilon = \left(\frac{M_p}{M_*}\right)\left(\frac{a_p}{\mathrm{AU}}\right)\left(\frac{d}{\mathrm{pc}}\right)^{-1} \quad \mathrm{arcsec} \tag{2}$$

where d is the distance from the observer to the planetary system.

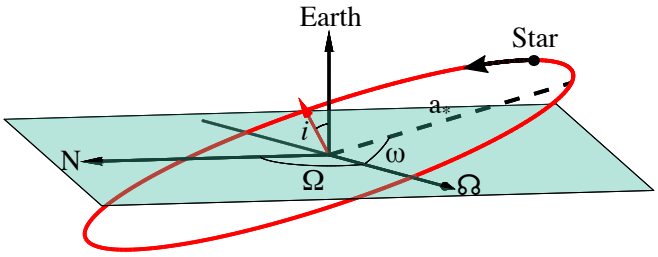

Fig. 1. Illustration of the four geometrical elements: the semimajor axis a_*, the longitude of the ascending node Ω, the inclination i, and the argument of periapsis ω. The light cyan polygon represents the sky plane in a perspective view; the reference direction is taken towards the North (N). The line of nodes, shown by the black arrow pointing towards the ascending node, is the intersection between the sky and the orbit planes. The orbit is shown as the red ellipse. The small red arrow is perpendicular to the orbit plane and shows the inclination.

The orbital elements define the orbit on the true plane. Three of them are called dynamical elements, as they define the motion in the orbit: the period P, the time of periastron passage T_p, and the eccentricity e. Four geometrical elements describe the size of the orbit and its orientation with respect to the plane of the sky: the semimajor axis a_*, the longitude of the ascending node Ω, the argument of periapsis ω, and the inclination i; see Fig. 1.

The instantaneous position on the sky of a star at time t is determined by: a fixed displacement $(\Delta\alpha_0, \Delta\delta_0)$ from its nominal position (α_0, δ_0) at a reference time t_0, the proper motion vector[2] (μ_α, μ_δ), the parallax ϖ, and the orbital elements. In Cartesian coordinates lying on a plane tangent to the sky and directed along right ascension and declination, respectively (Michalik et al. 2014), the position may be written as

$$\Delta\alpha(t) = \Delta\alpha_0 + \Pi_\alpha \varpi + (t - t_0)\mu_a + BX(t) + GY(t) \tag{3}$$
$$\Delta\delta(t) = \Delta\delta_0 + \Pi_\delta \varpi + (t - t_0)\mu_\delta + AX(t) + FY(t) \tag{4}$$

where Π_α and Π_δ are the components of the parallax direction, X and Y are the displacement due to planetary motion in the true plane, projected onto the sky through the Thiele-Innes (TI) constants (also called natural elements) A, B, F, G

$$A = \upsilon\,(\ \ \cos\omega\cos\Omega - \sin\omega\sin\Omega\cos i) \tag{5}$$
$$B = \upsilon\,(\ \ \cos\omega\sin\Omega + \sin\omega\cos\Omega\cos i) \tag{6}$$
$$F = \upsilon\,(-\sin\omega\cos\Omega - \cos\omega\sin\Omega\cos i) \tag{7}$$
$$G = \upsilon\,(-\sin\omega\sin\Omega + \cos\omega\cos\Omega\cos i) \tag{8}$$

[1] From the Greek $\upsilon\pi\sigma\gamma\rho\alpha\phi\acute{\eta}$ (signature).

[2] The proper motion component in the right ascension includes the $\cos\delta$ factor, and corresponds to the μ_{α^*} in Appendix A of Michalik et al. (2014).

In the true plane, the coordinates X and Y depend on the eccentric anomaly E and on the eccentricity e

$$X(t) = \cos E(t) - e \tag{9}$$

$$Y(t) = \sqrt{1-e^2}\,\sin E(t) \quad . \tag{10}$$

The eccentric anomaly is obtained from Kepler's equation

$$E(t) = M(t) + e\sin E(t) \tag{11}$$

where M is the mean anomaly

$$M(t) = \frac{360^\circ}{P}(t - T_p) \quad . \tag{12}$$

Within the TI formalism, the determination of the orbital elements from astrometric observations is a non-linear problem in P, e, and T_p, and linear in all other variables.

2.2. Degeneracies

In some cases, some parameters can become degenerate. For instance, in the case of circular orbits, the argument of periapsis ω loses its geometrical meaning as direction of the semi-major axis of the orbital ellipse, and it becomes degenerate with the periapsis transit time T_p.

A second case is that of face-on orbits, where $\cos i = \pm 1$. Here, the orbital plane coincides with the sky plane, so that both the longitude of the ascending node Ω and the argument of periapsis ω lie on the same plane and hence become degenerate.

A third case happens because astrometric data cannot distinguish the ascending node from the descending one. In principle, both the longitude of the ascending node Ω and the argument of periapsis ω vary between -180° and 180°. However, the same projected orbit can be obtained by adding 180° to both Ω and ω, i.e. by exchanging the ascending node for the descending node and the periapsis for the apoapsis. It is customary when analysing astrometric data to restrict the allowed range of Ω to the $[0^\circ, 180^\circ]$ interval (Binnendijk 1960; Perryman et al. 2014). This degeneracy could be broken if one were to jointly analyse astrometric and radial-velocity data.

The above degeneracies are observed in our analysis and the affected systems will be identified in Sect. 6.

3. Detectability as a function of the signal/noise ratio

Previous works by C08 and P14 have shown that the main parameters, upon which the detectability of a planet depends, are the period P and the signal/noise ratio SN

$$\mathrm{SN} = \frac{\upsilon}{\sigma_{\mathrm{AL}}} \tag{13}$$

where υ is the astrometric signature and σ_{AL} is the single-scan uncertainty on our observable, the shift in the along-scan (AL) field angle η (defined in Sect. 4.1).

Detectability may also depend on the period because orbit sampling becomes difficult or incomplete either if P is shorter than the the average time between two Gaia scans, or if P is longer than the mission length. This will be explored in Sect. 6.3.

Any value of SN can be mapped, using Eq. (2), to a subspace of M_p, M_*, a_p, d, and of the stellar magnitude V, which would produce the same simulated data. This is illustrated for a few example cases in Fig. 2, where a Neptune- and a Jupiter-sized planets are assumed to orbit a bright star at different semimajor axes and distances. Both C08 and P14 suggested that one needs SN $\gtrsim 3$ to detect a planet; therefore Fig. 2 also shows a first-order estimate of the distance from Earth up to which a planet can be detected.

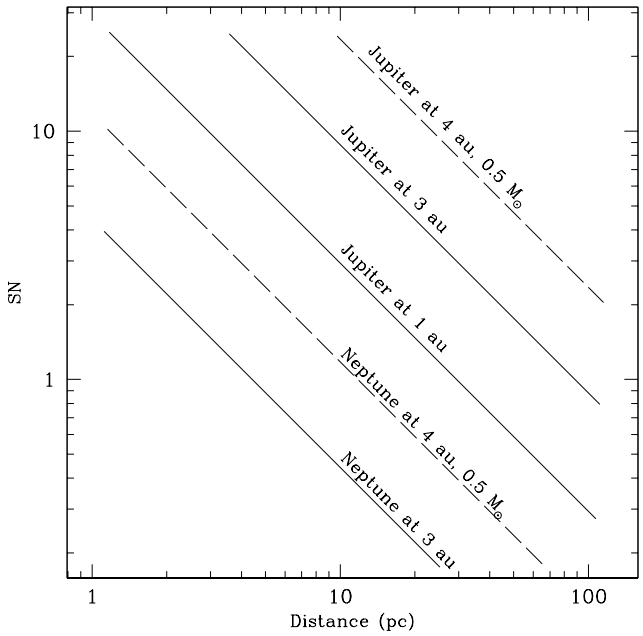

Fig. 2. Expected signal/noise ratio (SN) for different planet masses as a function of distance. The lines show the cases for 1 Jupiter mass and for 1 Neptune mass; with circular orbits at 1, 3, and 4 au (it is only possible to observe a 4 au orbit if the mission length is $\sim$ 10 yr; see Sect. 6.4). A $1\,M_\odot$ star is considered, except for the dashed lines where we show the case for $0.5\,M_\odot$.

Therefore, in the following we will parameterize our simulations over a grid of P and SN. Only when relevant will we trace the SN back to the parameters on which it depends.

4. Simulated Gaia observations

4.1. Field angles in the Gaia field of view

The Gaia spacecraft scans the sky along great circles whose orientation changes in time, following a complex pattern which we refer to as the Gaia scanning law. For each star transiting any of the two fields of view, the (time-dependent) position in the AL direction is measured on the sky mapper charge-coupled devices (CCD) and on each of the nine astrometric field CCDs. The formal uncertainty σ_{AL} results from combining all ten measurements, and it depends on the stellar magnitude, with $\sigma_{\mathrm{AL}} = 0.034$ mas for $V \leq 14$ for a single scan, and larger for fainter stars (de Bruijne 2012; de Bruijne et al. 2014). A coarser determination of the position in the across-scan (AC) direction is also done, but issues of pixel size and binning make AC measurements less useful for astrometry. Gaia returns on average 70 times to each star during its 5 yr mission (de Bruijne et al. 2010), and each time the transit in the field

of view measurement will occur with a different orientation, thus the stellar position can be pinpointed with AL information only.

If $\theta(t)$ is the position angle of the Gaia field of view, the shift in the AL field angle η is given by

$$\Delta\eta(t) = \Delta\alpha \sin\theta(t) + \Delta\delta \cos\theta(t) \tag{14}$$

and Eqs. (3, 4) can be expressed in terms of AL displacement:

$$\begin{aligned}\Delta\eta(t) = {} & \Delta\alpha \cos\theta(t) + \Delta\delta \sin\theta(t) \\ & + \Pi_\alpha \varpi \cos\theta(t) + \Pi_\delta \varpi \sin\theta(t) \\ & + (t-t_0)\mu_\alpha \cos\theta(t) + (t-t_0)\mu_\delta \sin\theta(t) \\ & + (BX(t) + GY(t)) \cos\theta(t) \\ & + (AX(t) + FY(t)) \sin\theta(t)\end{aligned} \tag{15}$$

4.2. Single-planet systems: simulations over a grid of parameters

Ideally, we aim to test how detection methods perform over the possible range of variation of all parameters. However, the high dimensionality of the problem of determining astrometric and Keplerian parameters together prevents a full exploration of all parameters on a grid.

For our single-planet simulations, we selected three parameters to be explored on a grid:

- the signal noise ratio SN;
- the period P;
- the eccentricity e;

and all other parameters were randomised using the probability distributions listed in Table 1 and explained in Sect. 4.3.

We produced a first batch of simulations for a 5 yr mission (the nominal Gaia mission length), and a second batch for a 10 yr mission, to investigate the improvements in case the mission is extended. In both batches, our grids ran with logarithmic increments over the ranges SN = 0.2–10, P = 0.08–10 yr, and e = 0.01–0.97. To smooth the grid and make the plots more readable, we added small random amounts to the SN, P and e of each simulation. The upper bound to the eccentricity corresponds to an ellipse with a 1:4 ratio between the semi-minor and semi-major axis.

4.3. Randomised astrometric parameters

We set the RA, Dec coordinates of the host star to random values, uniformly distributed over the celestial sphere. The Gaia scanning law was computed for each position using the AGISLab software (Holl et al. 2012), which was developed as part of the Gaia astrometric global iterative solution (AGIS, Lindegren et al. 2012).

We set the astrometric parameters ($\Delta\alpha$, $\Delta\delta$, ϖ, μ_α, μ_δ) to random values, according to the probability distributions listed in Table 1. We assumed $\Delta\alpha$ and $\Delta\delta$ to be normally distributed, with a standard deviation of 1 mas. The remaining parameters depend on the distance d, which was assumed to be log-normally distributed around a value which depends on the SN. The range of values covered by d was chosen so that the SN could include the cases of a Jupiter-mass planet at 3 AU around a $1M_\odot$ star, and of a Neptune-mass planet at 4 AU around a $0.5M_\odot$ star. The parallax was obtained from the distance. The proper motions should reflect the larger spread in velocities of nearby stars with respect to distant stars; therefore we drew proper motions from the Gaia data release 1 (DR1; Gaia Collaboration et al. 2016; Lindegren et al. 2016).

Table 1. Probability distributions of randomised parameters for the single-planet simulations.

Parameter	Distribution
RA, Dec	uniform over the celestial sphere
$\Delta\alpha_0$ (mas)	normal($\mu = 0$, $\sigma = 0.04$)
$\Delta\delta_0$ (mas)	normal($\mu = 0$, $\sigma = 0.04$)
d (pc)	log-normal($\ln\mu = \ln\frac{1\,\mathrm{mas}}{\sigma_{\mathrm{AL}}\ \mathrm{SN}}$, $\sigma = 0.6$)
ϖ (mas)	1000/d
μ_α, μ_δ (mas yr^{-1})	taken randomly from Gaia DR1 stars with distance $\sim d$
T_p (yr)	uniform(0, P)
Ω	uniform(0, 180°)
ω	uniform(−180°, 180°)
$\cos i$	uniform(−1, 1)

We set the remaining orbital elements (T_p, Ω, ω, i) to random values, different for each simulation, uniformly distributed over their natural ranges of variation: $0 \le T_p < P$, $0 \le \Omega < 180°$, $-180° \le \omega < 180°$, $-1 \le \cos i \le 1$.

Kepler's equation needs to be solved for every simulated Gaia scan; we used the numeric solver `vpasolve` in the MATLAB package, which provides solutions with double-precision accuracy[3].

4.4. Simulations of systems with no planet

To help setup the thresholds for planet detection, we also simulated no-planet systems. These systems follow the same specifications of the previous Sections for what pertains astrometric parameters, noise, and scanning law; but no Keplerian signal is added.

5. Bayesian analysis

In this Section, we assume some familiarity with Bayesian inference and its nomenclature (prior, likelihood, posterior, Markov chains, ...). The interested reader may find an introduction geared towards astrophysics in Trotta (2008); or a more general one in the textbook by Gregory (2005).

5.1. Modelling

From a model point of view, Bayesian inference works similarly to maximum likelihood, in that a series of tentative sets of parameters are generated, and each set is used to simulate an observable quantity which is compared against the data. In our setting, the data are $\eta_{0,s}$ i.e. the observed AL displacements, one per Gaia scan s, computed at simulation time using Eq. (15). Let $\eta_{i,s}$ be the displacements computed during the fit for a set i of astrometric and orbital

[3] We also experimented with the Mikkola (1987) cubic approximation, which has a relative error within 0.2%. We found that the use of the approximation was hindering the successful fit of the simulations with the method described in Sect. 5.

Table 2. Priors for the purely astrometric model. The parallax (ϖ) has the additional constraint that it must be non-negative. Proper motions have a parallax-dependent standard deviation.

Parameter	Distribution
$\Delta\alpha_0$ (mas)	normal($\mu = 0$, $\sigma = 0.04$)
$\Delta\delta_0$ (mas)	normal$(0, 0.04)$
ϖ (mas)	normal$(2.8, 7)$; $\varpi \geq 0$
μ_α (mas yr^{-1})	normal$(0, 9.5\,\varpi^{0.89})$
μ_δ (mas yr^{-1})	normal$(0, 10\,\varpi^{0.83})$

Table 3. Priors for the orbital elements of the planet. The full model includes both the orbital elements and the astrometric parameters of Table 2. The parameters for the uniform distribution are the lower and upper bounds of the allowed interval; the exponential distribution has the formula $P(x;\lambda) = \lambda e^{-\lambda x}$. The eccentricity is bound to the [0,0.97] interval whose upper bound corresponds to an ellipse with a 1:4 ratio between the semi-minor and semi-major axes.

Parameter	Distribution
P (yr)	log-normal($\ln\mu = -0.7, \ln\sigma = 0.25$)
υ (mas)	log-normal$(-2, 1)$
Ω	von Mises($\mu = 0, \kappa = 0.001$, support=$[0, 180°]$)
ω	von Mises(0, 0.001, $[-180°, 180°]$)
$\cos i$	uniform$(-1, 1)$
e	exponential($\lambda = 4.67$); $0 \leq e < 0.97$
T_p (yr)	von Mises(0, 0.001, $[0, P]$)

parameters. The likelihood may then be obtained by assuming that the difference $\eta_{0,s} - \eta_{i,s}$ is normally distributed, with mean zero and standard deviation $\sigma_{\rm AL}$:

$$\mathcal{L}(i) = \frac{1}{\sqrt{2\pi}\sigma_{\rm AL}} \prod_s e^{-\frac{1}{2}\frac{(\eta_{0,s}-\eta_{i,s})^2}{\sigma_{\rm AL}}} \quad . \tag{16}$$

We used the `Stan` package for statistical inference (Carpenter et al. 2016, version 2.12) to model our simulated Gaia observations. The `Stan` package allows models to be specified in a programming language very similar to that of popular packages such as BUGS (Lunn et al. 2000) or JAGS (Plummer et al. 2003). All these packages implement general-purpose Markov chain Monte Carlo (MCMC) samplers together with a high-level programming language to specify the model.

Our implementation follows the formalism of Sect. 2.1. For each simulated system, we define two models:

- a purely astrometric model, which only contains $\Delta\alpha_0$, $\Delta\delta$, ϖ, μ_α, and μ_δ. This model can be obtained by setting $A = B = F = G = 0$ in Eq. (15). This model is the baseline against which the full model will be compared when evaluating a detection;
- a full model, containing all the parameters of Eq. (15).

Kepler's equation needs to be solved at each MCMC iteration, but `Stan` does not provide a solver in its current version; therefore we use the Markley (1995) approximation, which is fast to compute and provides, in our implementation, a relative error of the order of 10^{-14}, close to the limits of a double precision variable.

5.2. Model priors

Purely astrometric model

The prior probability distributions for the model parameters should be set according to the expected range of variation, and to the expected distribution for their values.

For the astrometric model we chose very weakly-informative priors which only set the physical scale of the parameters (Table 2). For $\Delta\alpha_0$ and $\Delta\delta$ we used normal distributions with $\mu = 0$ and $\sigma = 0.04$ mas, consistently with the expected perturbation induced by a Jovian planet. For the proper motions μ_α and μ_δ we used normal distributions with $\mu = 0$ and σ that depends on the parallax, to reproduce the fact that more distant stars have lower velocities. By fitting the Gaia DR1 stars within ~ 200 pc, we found $\sigma_{\alpha_0} \propto \varpi^{0.83}$ and $\sigma_\delta \propto \varpi^{0.89}$. For the parallax ϖ, we used a log-normal prior (with $\ln\mu = 2.8$, corresponding to 16 mas, and $\ln\sigma = 7$) which we found to reasonably reproduce the Hipparcos and Gaia DR1 data. We also added the constraint that $\varpi \geq 0$.

Full model

For some parameters (Ω, ω, $\cos i$, and T_p) the obvious choice is to set the priors to constant distributions according to the parameter range; see Table 3. A uniform distribution is therefore our choice for $\cos i$. However, Ω and ω are angles, and T_p may also be treated as such, thus one needs to ensure that the prior distribution is defined over a circular support, not to introduce any artificial barrier for when the Markov chains explore the likelihood space. This is achieved by a von Mises distribution with concentration parameter $\kappa \sim 0$, that is the circular equivalent of a uniform distribution[4] (Jammalamadaka & SenGupta 2001).

For the other orbital parameters we chose weakly-informed priors as follows:

- For P, one should concentrate the probability on periods longer than the time between two Gaia scans, and shorter than the mission length.
  Therefore we chose a log-normal distribution with $\ln\mu = -0.7$ (corresponding to 0.5 yr) and $\ln\sigma = 0.25$, and we also restricted the allowed range to the 0.01–15 yr interval.
- For e, we considered the planets listed in the `exoplanet.eu` database, and observed that they have mostly small eccentricities. We found that an exponential distribution with a scale of 4.67 fits reasonably the eccentricities listed in the database, and ensures that $e < 0.5$ contains 90% of the probability.
- For the angular size of the semi-major axis υ we chose a log-normal distribution with $\ln\mu = -2$ (corresponding to 0.14 mas) and $\ln\sigma = 1$, with an upper bound at 10 mas. Such a distribution puts 90% of the probability on systems with $\upsilon < 0.5$ mas, yet it allows for some rarer systems closer to Earth where wider angular orbits might be expected.

Finally, astrometric parameters in the full model have the same priors as in the purely astrometric model.

[4] Ideally one should set $\kappa = 0$; however the current Stan implementation of the von Mises distribution is restricted to $\kappa > 0$, so we chose a κ low enough to be practically indistinguishable from 0.

5.3. Chain initialisation

Markov chains (MC) need a starting point. From there, the MC will start to explore the parameter space. If the starting point is not in a high-probability region, it is necessary to discard the first iterations (usually called burn-in period, or warm-up period) to ensure that the rest of the chain represents a sample of the posterior distribution, with no bias due to the presence of many initial steps exploring a low-probability region.

In principle, the starting point may be taken randomly, though this may result in a very long burn-in period. Therefore, statistical literature often suggests to start the Markov chains at a point reasonably close to the likelihood global maximum, if this is known (e.g., Geyer 2011).

Our tests have shown that the MC burn-in can be prohibitively long unless a starting point close to the correct values for the parameters is used. However, in our case the structure of the likelihood space presents many very narrow peaks, that prevent common optimisation methods (e.g. Amoeba, Levenberg-Marquardt, or Broyden–Fletcher–Goldfarb–Shanno) from finding the global likelihood maximum. To find a suitable starting point, we have considered a simplified model with circular orbits in which the period P is the only parameter which appears non-linearly. We defined a grid of frequencies $f = 1/P$, starting from $f = 0.067$ to 12.5 yr^{-1}, with a step of 0.03 yr^{-1}, covering the allowed interval for P. For every value of f, we obtained the best-fit values for all other parameters and the χ^2 using linear least-squares. The values which scored the lowest χ^2 provided the starting point for the MC.

The circular-orbit model only includes 10 parameters: e does not appear; and ω and T_p become degenerate, so that only one of them is needed. To make our implementation easier, we dropped T_p and retained ω. Therefore, we were left with the need to specify starting values for e and T_p. We chose $e = 0.1$ and $T_p = 0$.

5.4. Chain convergence

In the long run, the set of points touched by a MC forms a non-biased sample of the target probability distribution; when this happens, it is said that the MC has converged to the target. The main risk associated with non-converging chains is that, since the MCMC samples would not be representative of the posterior distribution, any inference on the model parameters might be severely biased. The purpose of the burn-in period is to assure that the chains have converged before one starts getting posterior samples from them. However there is little guidance on how long the burn-in period should be, so one has to rely on experience and on inspection of the MCMC samples.

We initially conduced some test runs of our `Stan` implementation to understand the number of iterations needed for the Markov chains burn-in period. Visual inspection of the chain traces showed that the burn-in period lasted less than a few hundred iterations in many cases where SN $\gtrsim$3–5; but in some cases it could be considerably longer.

Therefore, for the simulations described in Sect. 4.2, we initially decided to use a burn-in period of 500 iterations. Then we would check if the chain had converged after this period, or if more burn-in iterations were needed, according to the procedure described as follows.

After the 500 burn-in iterations, we collected the next 500 iterations for further analysis. On the latter samples we computed for each parameter i the potential scale reduction $\hat{R}_i$, that measures how much sharper an estimate might become if the MCs were run indefinitely (Gelman & Rubin 1992; Eq.(11.4) in Gelman et al. 2013). The $\hat{R}$ compares the variance among the chains vs. the variance within a single chain: if all chains have converged and are sampling the same space, both variances should be approximately equal, and $\hat{R} \sim 1$. A value of $\hat{R} \gtrsim 1.1$ is usually considered as pointing towards convergence problems (though the opposite is not true: it is possible to have $\hat{R} < 1.1$ and still have convergence problems).

We found that, for any given simulated system, the $\hat{R}_i$ can vary among the parameters, with some of them showing convergence problems while other seem to be fine. Therefore we also computed the average $\hat{R}_{\mathrm{avg}} = \frac{1}{N}\sum_{i=1}^{N}\hat{R}_i$. Whenever a simulated system was found to have $\hat{R}_{\mathrm{avg}} > 1.1$ we discarded its samples and we ran the chains anew, this time with 1500 burn-in iterations. Then, 500 iterations were collected for further analysis, and the $\hat{R}$'s were re-assessed. If the $\hat{R}_{\mathrm{avg}}$ criterion was still not satisfied, we repeated the procedure with 4500 or, if needed, 13500 burn-in iterations.

After that, if $\hat{R}_{\mathrm{avg}}$ was still unsatisfying, we deemed the chains as not converging and excluded the planet from further consideration, except for a short discussion in Sect. 7.3.

There may be different reasons for not fulfilling the convergence criterion, some of which do not prevent meaningful information to be retrieved. It would be however difficult to express how confident one could be in any estimate coming from non-converging chains. Therefore, in the context of a blind survey of exoplanets, information from non-converging chains might be of little use; and in the context of extracting most information from Gaia data of a particular planet (say, by joint fitting of astrometric and radial velocity, or transit data) there would be probably more prior information to help find a different (perhaps better) starting point for the chain.

5.5. Exoplanet detection as model selection

Deciding whether an exoplanet has been detected is essentially a comparison of model performances. The null model is that no planet is present, and that the Gaia measurements may be adequately explained with just the astrometric parameters. The alternative model is that a planet is present, and both the astrometric and Keplerian parameters are needed.

The classical (frequentist) approach is to use hypothesis testing: the difference between the models' χ^2 is computed and linked, through the number of degrees of freedom, to a p-value. It can be applied to the case of exoplanet detection, provided that the Thiele-Innes parameterization is adopted. One would therefore obtain the A, B, F, and G parameters from the fit, and only later they would be reconnected to the orbital elements υ, ω, Ω, and $\cos i$. Uncertainties on the best-fit values should then be propagated from the Thiele-Innes parameters back to the orbital elements.

However, a complication is that the number of degrees of freedom is unambiguously defined only for linear models. Non-linear models behave like they have a larger number of degrees of freedom than a linear model with the same num-

ber of parameters. Also, the "effective degrees of freedom" of a non-linear model are data-dependent (through the variance of the parameters; see e.g. Eq. (19) below), and they are larger when data are noisy (Andrae et al. 2010); we will show this happening in our simulations in Sect. 6.1.

Conversely, the Bayesian approach has the following advantages: the orbital elements can be the subject of inference, the uncertainties on the fit parameters need not to be estimated using asymptotic theory, and there exist model comparison methods which explicitly account for the effective degrees of freedom.

An easy way for Bayesian model comparison is to use information criteria (IC), because only the posterior samples are needed for their calculation. The theory of most IC is based on the notion of predictive accuracy, i.e. how well a model fit can predict new data produced from the true data-generating process. A notable exception is the Bayesian information criterion (BIC, see below) which is motivated by the marginal probability of the data under the model. A introduction to both criteria from an astrophysics perspective can be found in Liddle (2007); among statistical sources, we refer to Gelman et al. (2013, Chap. 7) and Burnham & Anderson (2002).

The IC operate in practice by computing a function of the likelihood, either at a specified point or averaged over the posterior, and then adding a penalisation for the number of degrees of freedom.

The best known IC are the Akaike information criterion (AIC, Akaike 1973)

$$\mathrm{AIC} = -2\ln\mathcal{L}_{\max} + 2k \tag{17}$$

and the Bayesian information criterion (BIC, Schwarz 1978)

$$\mathrm{BIC} = -2\ln\mathcal{L}_{\max} + k\ln n \tag{18}$$

where $\mathcal{L}_{\max}$ is the maximum likelihood, k is the number of free parameters, and n is the number of data points. The main difference is that the penalty on BIC also depends on the number of data points.

Both the AIC and the BIC rely on a point estimate of the model, and they still involve counting the free parameters. A recent development on the AIC is the widely applicable (or Watanabe-Akaike) information criterion (WAIC), that has been proposed as a fully Bayesian approach of estimating the predictive accuracy (Watanabe 2010; see Ranalli et al. 2016 for an application to astrophysics). It uses all available samples of the posterior distribution and gets the effective degrees of freedom from the parameter variance

$$\mathrm{WAIC} = -2\ln\left(\frac{1}{I}\sum_{i=1}^{I}\mathcal{L}(i)\right) + 2\sum_{s=1}^{n}\frac{1}{I-1}\sum_{i=1}^{I}\left(\ln\mathcal{L}(i,s) - \langle\ln\mathcal{L}(i,s)\rangle_i\right)^2. \tag{19}$$

Here, the index i runs on the samples from the posterior distribution, and the index s runs on the data points (the Gaia scans), so both indices keep the same meaning they had in Eq. (16). The symbol $\mathcal{L}(i,s)$ is the likelihood of a single data point (scan). The angular brackets with the i subscript indicate that the average over the samples should be taken: the squared difference is therefore the variance of the log-likelihood[5].

Table 4. Information criteria and $\Delta\chi^2$ thresholds for 99.7% (3σ) detection of planets; and their maximum values observed in no-planet simulations.

Criterion	3σ threshold		Maximum	
	(5 yr)	(10 yr)	(5 yr)	(10 yr)
WAIC	16.8	19.3	24.9	26.5
AIC	13.2	17.4	23.6	21.9
BIC	-2.78	-2.55	3.99	2.24
$\Delta\chi^2$	26.8	31.6	37.6	35.9

In all cases, the lower the IC, the more favoured the model. However, the absolute scale of the IC does not carry statistical meaning; the important quantity is the difference between the IC of different models, not the values of the IC themselves. Therefore, a large ΔIC can indicate a strong preference for one model over the other, while a small ΔIC indicates that one model does not perform appreciably better than the other.

The preference can be quantified by considering the relative likelihood w of model M_1 over model M_2. Assuming $\mathrm{IC}(M_1) > \mathrm{IC}(M_2)$ (so that M_2 is the preferred model),

$$w = e^{-\frac{1}{2}\Delta\mathrm{IC}_{12}} \tag{20}$$

where $\Delta\mathrm{IC}_{12} = \mathrm{IC}(M_1) - \mathrm{IC}(M_2)$ (Burnham & Anderson 2002, see their Sects. 2.8 and 2.9). However, the relative likelihood w needs some care in its interpretation, for different w are obtained when different IC are applied to the same data. This has to do with the definition of each IC and with the fact that AIC/WAIC and BIC are derived under very different hypotheses; we prefer to refer to the statistical literature for further discussion (Burnham & Anderson 2002, see their Chap. 6).

The ΔIC can also be interpreted using Jeffreys' scale of evidence (Efron & Gous 2001), which rates $\Delta\mathrm{IC} \gtrsim 1$, 3, 20, and 150 as "barely worth reporting", "weak", "strong", and "very strong", respectively. Strong evidence would correspond to a relative likelihood of $w \sim 4.5 \times 10^{-5}$ for the worse model over the preferred one.

In Sect. 6 we will empirically calibrate our ΔIC thresholds for planet detection.

6. Results

In this Sect., we discuss results from the Bayesian fits described in Sect. 5, and we only consider systems for which the MC convergence criterion defined in Sect. 5.4 was satisfied. Non-converging systems will be discussed in Sect. 7.3.

6.1. Detection threshold

6.1.1. Information criteria

In this Sect. we apply the three criteria to data where we expect that the purely astrometric model should be pre-

[5] We note that Watanabe (2013) also proposed an extension of the BIC: the widely applicable Bayesian information criterion (WBIC). However, to compute the WBIC one needs samples from both the posterior and the prior. The latter are not provided by the current Stan version, and therefore we do not consider the WBIC here.

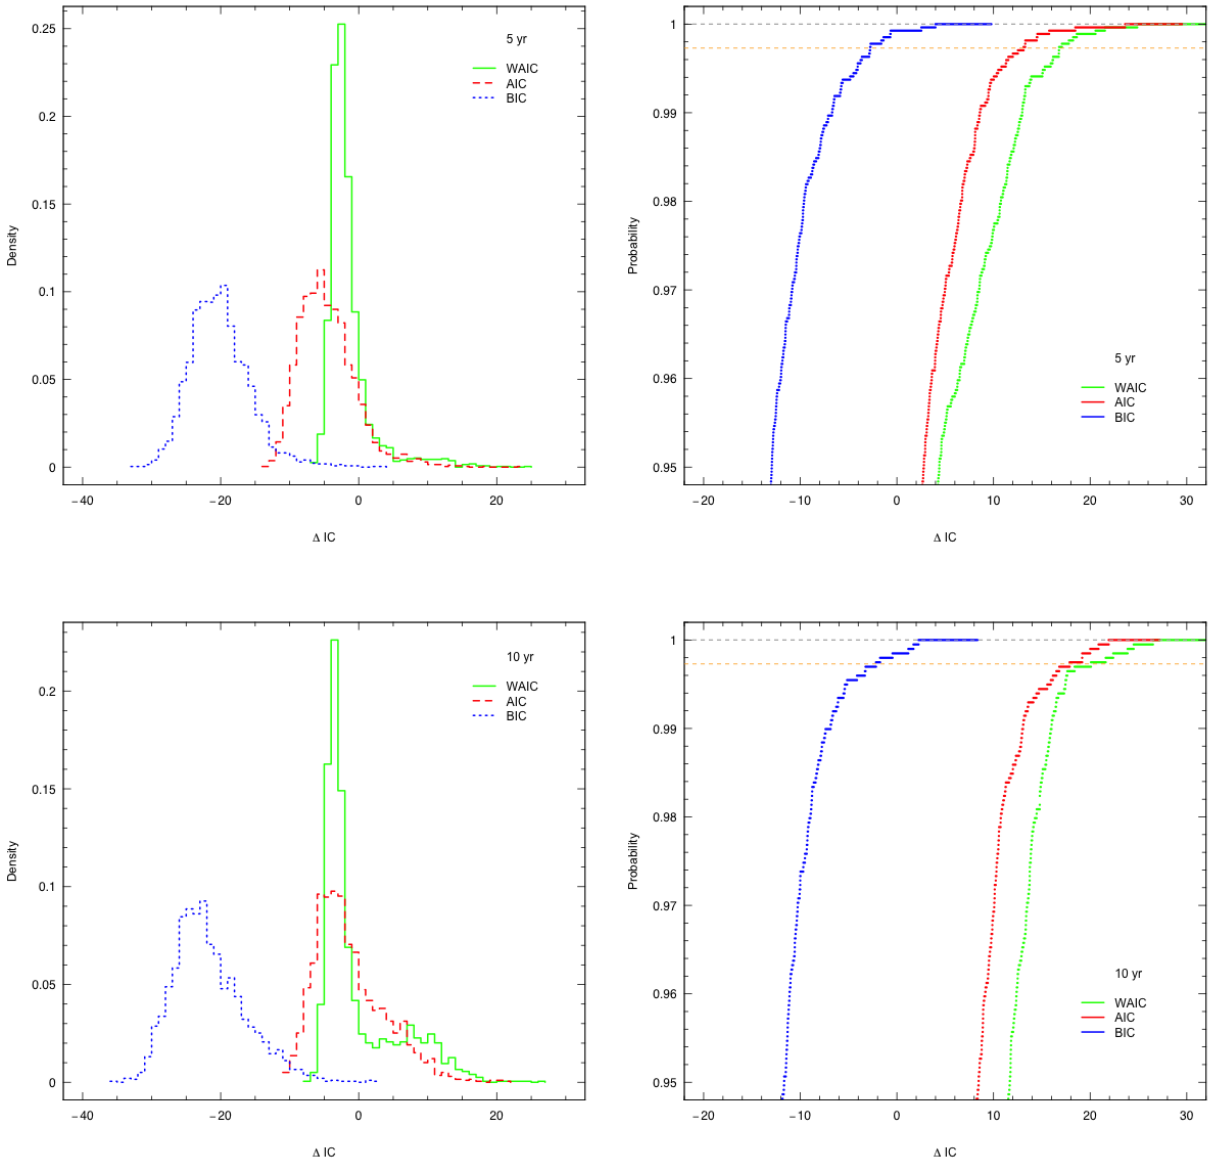

Fig. 3. Histograms (left panels) and cumulative distribution (right panels) for the three information criteria in the case of no-planet simulations. The upper panels refer the 5 yr mission, and the lower panels to the 10 yr mission. Left panels: green solid lines, WAIC; red dashed lines, AIC; blue dotted lines, BIC. Right panels: from left to right, BIC, AIC, and WAIC are shown with blue, red, and green lines, respectively; the horizontal orange dashed line shows the 99.73% (3σ) quantile.

ferred, i.e. our no-planet simulations. Our samples consist of 2712 and 1986 no-planet systems for the 5 yr and 10 yr missions, respectively.

We define $\Delta\mathrm{IC} = \mathrm{IC}(M_a) - \mathrm{IC}(M_k)$, where M_a is the astrometric model, and M_k is the model with the Keplerian orbit. With our definition, a positive ΔIC indicates a preference for the Keplerian model, and vice versa. Values of $\Delta\mathrm{IC} \sim 0$ indicate no preference for a model over the other; therefore the simpler model (i.e. the astrometric model) should be preferred.

In Fig. 3, we show the histograms and the cumulative distributions of ΔIC for the no-planet systems. It can be seen that ΔWAIC and ΔAIC are mostly contained below a value of 5, but with maximum observed values in the 20–30 interval; we report the 99.7% quantiles and the maxima in Table 4. When going from the 5 yr to the 10 yr mission, there is a slight increase in the 99.7% quantile for all IC. The maximum ΔWAIC also increases, but the maximum ΔAIC and ΔBIC decrease: this is probably due to the maximum being a less stable quantity than a quantile. Assuming a $\Delta\mathrm{IC} \sim 20$ as an approximate threshold for strong evidence according to Jeffreys' scale, Table 4 shows that for more than 99.7% of simulated no-planet systems, the evidence for the Keplerian model is less than strong. Since the ΔAIC and ΔWAIC maxima occur at values in the 20–30 range, we can conclude that the tails of ΔIC does not extend much above 20.

Compared to ΔAIC and ΔBIC, ΔBIC appears to be much more conservative, almost always having negative values and therefore indicating a preference for the astrometric model (ΔBIC was positive in 2 cases out of 2712 for the 5 yr mission, and in 3 out of 1986 for the 10 yr mission).

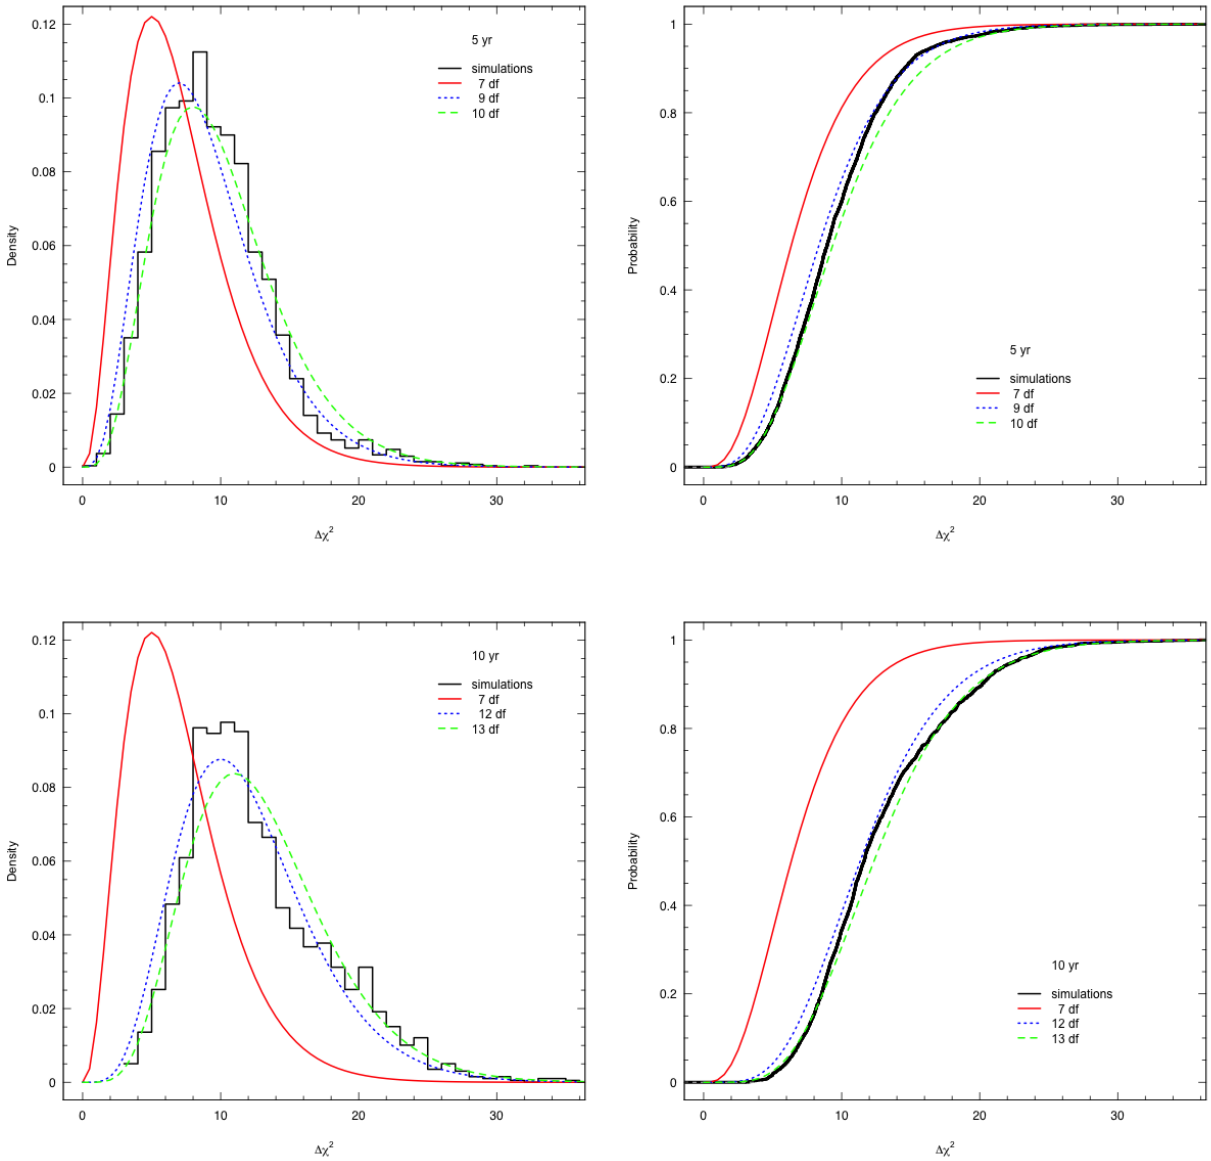

Fig. 4. Histograms (left panels) and cumulative distributions (right panels) of $\Delta\chi^2$ for no-planet systems in the 5 yr (upper panels) and 10 yr (lower panels) mission simulations, showing the empirical density of $\Delta\chi^2$ for systems with no detectable planet. The solid and dashed curves show χ^2 distributions with the degrees of freedom (dof) reported in the legend. The empirical density does not follow a χ^2 distribution with the nominal number of dof (7, shown as the solid red line). It can however be approximated, especially in its right tail, by a χ^2 distribution with a larger number of dof: 9 dof for the 5 yr mission, 13 dof for 10 yr mission.

Setting a threshold for planet detection is therefore a somewhat arbitrary operation. Depending on one's tolerance for false positives and on the size of the sample of stars to be analysed, the threshold could be as low as 20, or much higher. An argument for a higher threshold could be that in systems with a planet that has a low SN, the recovered parameters could have larger associated uncertainties; thus, one could prefer a higher threshold if that allowed to select systems with higher SN. This argument is the subject of Sect. 6.2, where we will discuss again the threshold.

The efficiency of each criterion (included the $\Delta\chi^2$, discussed below) in detecting planets also depends on the chosen threshold. In Sects. 6.2 and 6.3 we will use a common thresholds for all criteria; this however makes some criteria more efficient than the others. Therefore, in Sect. 7.4 we will return on the subject, to show that by tuning the thresholds all criteria can achieve the same efficiency.

6.1.2. $\Delta\chi^2$

Methods based on the χ^2 are widely used for hypothesis testing, therefore it may be interesting to compare against the IC.

Considering two nested models with a different number of parameters (for instance, the astrometric and the full model), and data distributed according to the null model (in our case, the astrometric model), the difference between the χ^2 values of the best fits under the two model models still follows a χ^2 distribution, with a number of degrees of freedom (dof) equal to the difference between the dof of the

two models. As discussed in Sect. 5.5, dof become ill-defined when the model is non-linear. We show this in Fig. 4, where we plot the empirical $\Delta\chi^2$ density for the no-planet simulations, against a χ^2 distribution for 7 dof (corresponding to the difference between 12 and 5 parameters in the full and astrometric models, respectively). The empirical density is better described by a χ^2 distribution with 9 or 10 dof for the 5 yr mission (the best approximation is obtained with 10 dof in the left tail, and 9 dof in the right tail); and with 13 dof for the 10 yr mission (with some deviations around the peak).

From the empirical distribution one can obtain a 99.7% quantile and a maximum, as in the case of the information criteria; these values are reported in Table 4. As in the case of ΔAIC and ΔBIC, the quantile increases when going from the 5 yr to the 10 yr mission, while the maximum decreases. If one were to use $\Delta\chi^2$ to set a detection threshold, it should be at least larger than the 99.7% quantile ($\Delta\chi^2 \sim 26.8$–31.6), and possibly even more conservative. Some guidance could be offered by the χ^2 distributions with 9 dof and 13 dof discussed above; or by an additional run of no-planet simulations with a larger sample size, so that the right tail of the distribution could be populated enough to allow the determination of a 4σ- or 5σ-equivalent quantile.

6.2. Detection metrics and signal/noise ratio

6.2.1. Information criteria

In Fig. 5 we compare the ΔAIC, ΔBIC, and ΔWAIC for our simulations against the SN, for a 5 yr mission (left panel) and a 10 yr mission (right panel). The AIC and WAIC behave rather similarly, while BIC is more conservative, as already seen.

The transition between non-detections and detections depends on the mission length. For the 5 yr mission, the transition starts around SN $\sim$ 1 for ΔWAIC and ΔAIC, and the running medians of both ΔWAIC and ΔAIC cross the ΔIC = 20 threshold (strong evidence in Jeffreys' scale) at SN $\sim$ 1.5. The ΔBIC starts the transition at SN $\sim$ 1.2 and crosses ΔIC = 20 at SN $\sim$ 1.9.

For the 10 yr mission, lower SN are needed for the detection. The transition starts around SN $\sim$ 0.7 for ΔWAIC and ΔAIC, and the running medians of both ΔWAIC and ΔAIC cross the ΔIC = 20 threshold at SN $\sim$ 1. The ΔBIC starts the transition at SN $\sim$ 0.9 and crosses ΔIC = 20 at SN $\sim$ 1.3.

From Fig. 5 it appears that there is no self-evident value to set a ΔIC threshold for candidate detections. On the one hand, a lower threshold assembles a larger number of detections; on the other hand, a high threshold should improve the quality of the detection, in the sense of not including systems whose parameter estimates are noisier or reflect secondary likelihood maxima. In the region where $1.5 \lesssim$ SN $\lesssim 1.9$ (for the 5 yr mission) and $0.7 \lesssim$ SN $\lesssim 1.3$ (for the 10 yr mission), the AIC and the WAIC signal a detection in the majority cases, while the BIC signals a non-detection. Our tests showed that in these regions, and also for higher SN up to ~ 2.3 (5 yr) and ~ 1.5 (10 yr), the best-fit parameters were incorrectly recovered in a number of cases (e.g., for the period it happened in approximately one quarter of the cases). Therefore, we prefer a more conservative threshold, that we set at ΔWAIC= 70, and which is roughly equivalent to ΔAIC=70 and ΔBIC= 20. We will use this threshold in the remaining of this paper.

6.2.2. $\Delta\chi^2$

The $\Delta\chi^2$ is also correlated with the SN, as shown in Fig. 4. Following the discussion in the previous Sect., we mark the $\Delta\chi^2$ thresholds corresponding to the 99.7% and 99.99994% quantiles for no-planet systems (3σ and 5σ, respectively); the lower quantile is obtained from the empirical cumulative distribution, while the upper quantile is calculated from the χ^2 approximations with 9 dof and 13 dof for the 5 yr and 10 yr mission, respectively.

Considering the 5σ threshold and the 5 yr mission, the transitions from non-detections to detections start around SN $\sim$ 1.1, while the running medians of $\Delta\chi^2$ cross the 5σ threshold around $SN \sim 1.8$. For the 10 yr mission, the transition starts around SN $\sim$ 1, while the running medians cross the threshold around $SN \sim 1.3$.

We conclude that the $\Delta\chi^2$ can also be used as a detection metric, provided that the number of degrees of freedom are empirically calibrated. Choosing a detection threshold is however no less arbitrary than in the case of information criteria, though some guidance is offered in both cases by well-established rules-of-thumb (Jeffreys' scale, vs. the 5σ rule). While one could relate a $\Delta\chi^2$ value to a ΔIC, such a relationship would only be valid for the problem at hand; and while at it, one should be careful not to mix Bayesian and frequentist statistics together, for they rely on different definitions of probability. For the remaining of this paper, we will not consider the $\Delta\chi^2$ further.

6.3. Detection rate

The detection efficiency of the three IC is presented in Fig. 7, where the fraction of detections is plotted for different SN. Using a threshold of ΔIC = 70, the detection efficiency is similar for the AIC and the WAIC, while the BIC detects 15%–30% less planets when $3 \lesssim$ SN $\lesssim 4$.

The period also affects the detectability in some cases. In Fig. 8 we show the detection fraction (using WAIC) as a function of both SN and period. The transition from non-detections to detections happens, as expected from the above discussion, at SN $\sim$ 2. However, the detections also exhibit a notch for systems with $2 \lesssim$ SN $\lesssim 4$ and P longer than the mission length.

6.4. Parameter recovery

In the rest of this Section, we will only consider candidate detections, i.e. the cases where ΔWAIC= 70 was obtained, and look at how well the orbital parameters are recovered. For every parameter, we calculated the median (the circular median is used for Ω, ω, and T_p; see Pewsey et al. 2013) and an estimate of the associated error: for Ω, ω, and T_p we use the circular standard deviation[6]; for every other parameter, we use the 68.3% highest posterior density interval (HPDI)

[6] Defined as $\hat{\sigma} = \sqrt{-2\ln\bar{R}}$, where $\bar{R} = \frac{1}{n}\sum_{i=1}^{n}\cos(\gamma_i - \bar{\gamma})$ is the sample mean resultant length, for a sample of n of angles γ_i whose mean is $\bar{\gamma}$; and calculated using the `sd.circular` function in the `circular` package for the R software; see Pewsey et al. (2013).

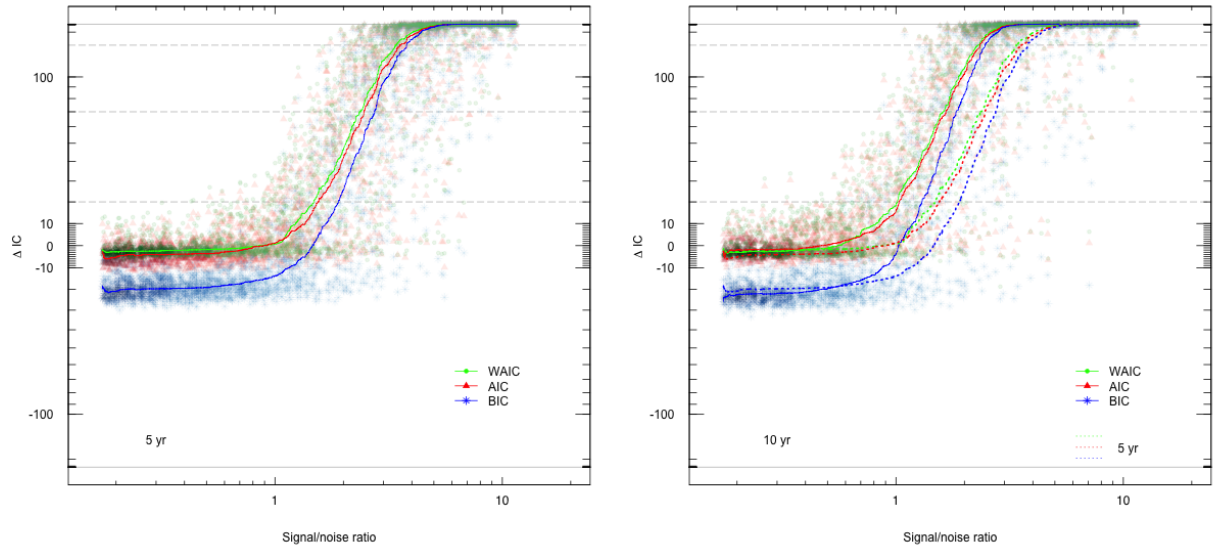

Fig. 5. Comparison of three different information criteria vs. signal noise ratio. Values of $\Delta\mathrm{IC} \sim 0$ indicate no preference for a model over the other; therefore the simpler model should be preferred. Negative and positive values of $\Delta\mathrm{IC}$ indicate a preference for the purely astrometric and full model, respectively. The horizontal grey dashed lines mark $\Delta\mathrm{IC} = 20$, 70, and 150; corresponding to strong evidence in Jeffreys' scale, our choice in this paper, and very strong evidence, respectively. Left panel: 5 yr mission; right panel: 10 yr mission. Green circles: WAIC; red triangles: AIC; blue stars: BIC. Darker colours indicate the presence of many overlapping points. Solid lines: running medians of the WAIC, AIC, and BIC. The solid lines of the left panel are repeated as dotted lines in the right panel, so that an immediate comparison between 5 yr and 10 yr is possible. The peculiar scaling on the y-axis is an inverse-logit transform ($f(\Delta\mathrm{IC}) = 1/(1 + e^{-0.02\Delta\mathrm{IC}})$) which keeps an approximately linear scaling around $\Delta\mathrm{IC} \sim 0$ and shrinks at large $|\Delta\mathrm{IC}|$, so that the $[-\infty, +\infty]$ range can be represented on the plot, and in fact any $|\Delta\mathrm{IC}| \gtrsim 300$ occupies approximately the same y positions on the plot, shown by the two horizontal grey solid lines).

from the posterior samples. Such quantities are plotted in Figs. 9–16.

Period

Fig. 9 shows that the periods of systems with $P_{\mathrm{input}} \lesssim T$, where T is mission length hence either 5 or 10 yr, are mostly correctly recovered.

A small number of systems ($\sim 8\%$) is fit with periods that are somewhat shorter or longer than the input one by a factor $\sim$ 40–60% (they are visible in Fig. 9 as the two parallel sequences right above and below the one-to-one line); these systems are also mis-fit as edge-on, high-eccentricity systems (they have incorrect $\cos i_{\mathrm{output}} \sim 0$ and $e_{\mathrm{output}} \sim 1$).

Systems with $P_{\mathrm{input}} \gtrsim 0.8T$ have a P_{output} that flattens at a value just below the length of the mission; this is also reflected in lower $\upsilon_{\mathrm{output}}$ than expected (see below).

Size of the orbit

The angular size of the semi-major axis of the orbit is given by the astrometric signature υ; its input value $\upsilon_{\mathrm{input}}$ can be obtained from the SN through Eq. (13). In Fig. 10, we show that the one-to-one relationship is well recovered.

Some systems deviate from the relationship, more so at larger υ; the deviations are present on both sides of the relationship for the 5 yr mission, but almost only on the right side for the 10 yr mission. Most of the systems with deviations on the right side (i.e. with $\upsilon_{\mathrm{output}} < \upsilon_{\mathrm{input}}$) have periods longer than the mission length, and are those that appear as the horizontal plume at $\sim 80\%$ of the mission length in Fig. 9; this behaviour is consistent with Kepler's law.

Conversely, most of the systems with deviations on the left side (i.e. with $\upsilon_{\mathrm{output}} > \upsilon_{\mathrm{input}}$) are also mis-fit as edge-on, high eccentricity systems ($\cos i_{\mathrm{output}} \sim 0$ and $e_{\mathrm{output}} \sim 1$).

The bottom left corner of the plot is more populated for the 10 yr mission, consistently with more planets being detected at low SN in the 10 yr mission and in the 5 yr one.

Eccentricity

The eccentricity (Fig. 11) exhibits quite a large scatter, both in terms of the location of the median e_{output} and of the size of the HPDI. The scatter is much reduced in the case of the extended mission, where the one-to-one relationship between e_{input} and e_{output} becomes much more evident. The relationship is also partly hidden by the fact that Fig. 11 is shown on a linear scale, while eccentricities were assigned in logarithmic scale, with more systems in near-circular orbits and in eccentric ones; this explains the concentration of points at $e_{\mathrm{input}} \lesssim 0.15$ and $e_{\mathrm{output}} \lesssim 0.15$.

The deviations from the one-to-one relationship seem to belong to two different groups. The first group includes systems with $e_{\mathrm{input}} \sim 0$ that are incorrectly fit with $e_{\mathrm{output}} \gg e_{\mathrm{input}}$; this happens in 41% (25%) of the systems with $e_{\mathrm{input}} < 0.15$ for the 5 yr (10 yr) mission. The second group includes systems with any e_{input} that are fit with $e_{\mathrm{output}} \sim 1$ with no dependence on e_{input}; this happens in

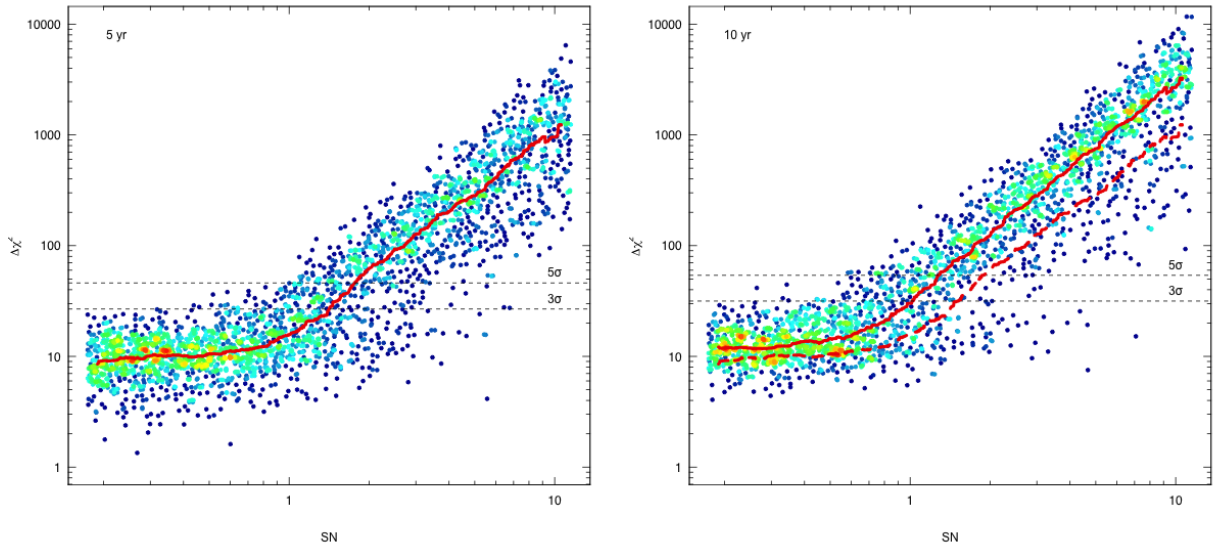

Fig. 6. Comparison of $\Delta\chi^2$ vs. SN, for the 5 yr mission (left panel) and the 10 yr mission (right panel). The colours show the density of points where they overlap, going from blue (no overlap) to green, yellow, and red (maximum overlap). The brown solid lines show the running medians of $\Delta\chi^2$; in the right panel, we also report the running median for the 5 yr mission as the brown dashed line. The horizontal grey dashed lines show the $\Delta\chi^2$ corresponding to the 99.7% and 99.99994% quantiles for no-planet systems (3σ and 5σ, respectively).

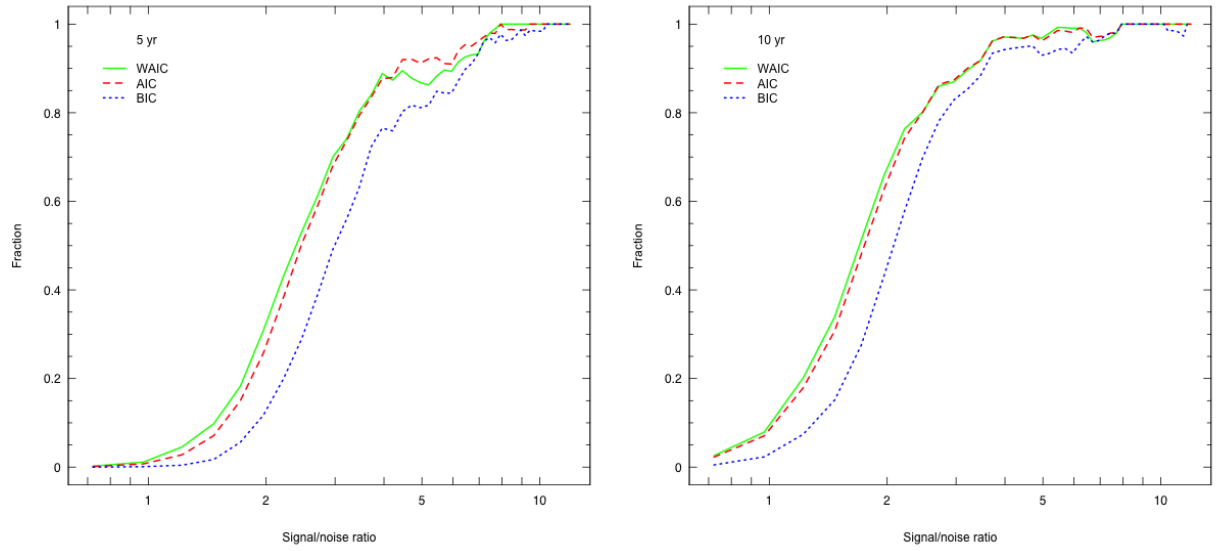

Fig. 7. Detection fractions under three different information criteria vs. signal noise ratio. Green solid line: WAIC; red dashed line: AIC; blue dotted line: BIC. Left panel: 5 yr mission; right panel: 10 yr mission.

17% (9%) of the systems with $e_{\rm input} < 0.8$ but $e_{\rm output} > 0.8$ for the 5 yr (10 yr) mission.

The systems in the second group also exhibit incorrect fits in other variables: they have $\omega_{\rm output} \sim \pm 90°$, an incorrect T_p, and $\cos i_{\rm output} \sim 0$ (edge-on); their $P_{\rm output}$ also run on the parallel sequences to the one-to-one relationships.

Some of the systems with $e_{\rm output} \sim 1$ and $\cos i_{\rm output} \sim 0$ are also associated with large deviations in $\Delta\alpha_{\rm output}$ and $\Delta\delta_{\rm output}$, though such associations are not one-to-one.

Inclination

The inclination is recovered correctly in the majority of cases (Fig. 12); the scatter on the one-to-one relationship is reduced when going from the nominal to the extended mission.

Most deviations from the one-to-one relationship seem to occur at moderate to high inclination ($0.5 \lesssim |\cos i| \lesssim 1$; $|\cos i| = 1$ corresponds to face-on systems); about half of the systems exhibiting such deviations also have $e_{\rm output} \sim 1$ and $\omega_{\rm output} \sim \pm 90°$, and an incorrect Tp.

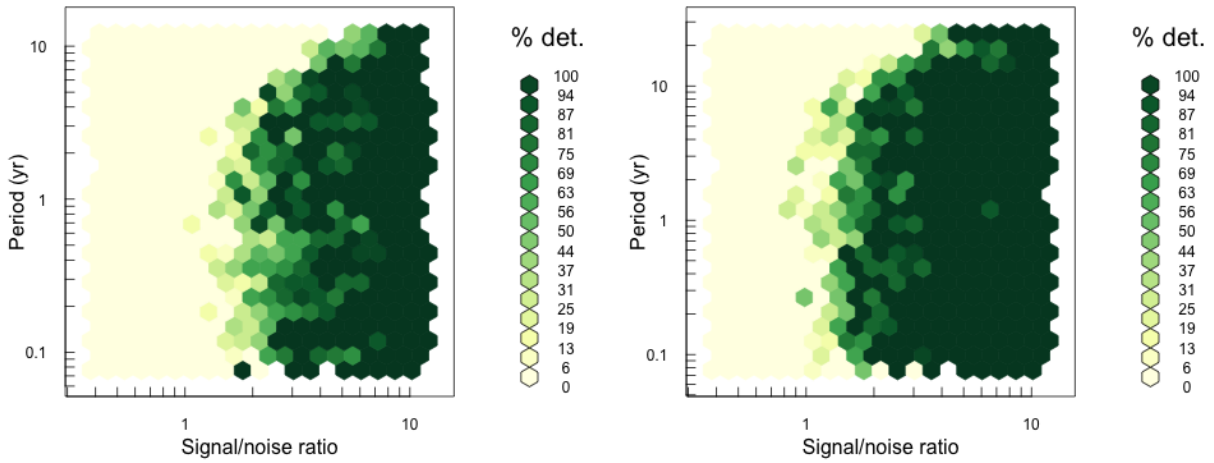

Fig. 8. Detection fractions as a function of SN and period. Left panel: 5 yr mission; right panel: 10 yr mission. The colours show the detection fraction and range from light yellow (low detection fraction) to dark green (high detection fraction). The transition between non-detections and detections is sharper and occurs at lower SN in the 10 yr mission than in the 5 yr mission. Planets with periods longer than the mission length require a stronger SN to be detected than planets with lower periods.

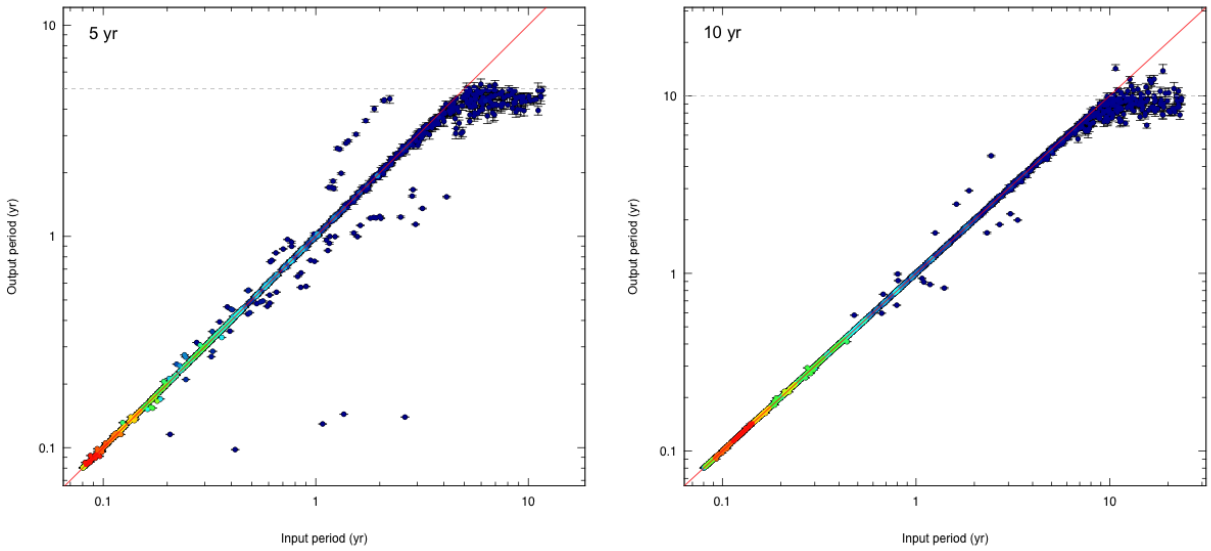

Fig. 9. Simulated vs. recovered period. The left and right panels show the 5 yr and 10 yr missions, respectively. The colours show the density of points where they overlap, going from blue (no overlap) to red (maximum overlap). The error bars show the 68.3% highest posterior density interval. The red solid line shows the one-to-one relationship. The period is correctly recovered for most of the simulated systems with periods shorter than the mission length. The two sequences parallel to the main relationship are populated with systems that are incorrectly fit with edge-on, high-eccentricity orbits; these parallel sequences are much less populated in the 10 yr mission.

Longitude of ascending node

The longitude of the ascending node (Ω) is correctly recovered in most cases; see Fig. 13. There seems not to be any particular structure in the points that deviate most from the one-to-one relationship.

Argument of periapsis

Whether the argument of periapsis (ω) is correctly recovered depends on the eccentricity. As discussed in Sect. 5.3, when orbits are nearly circular, ω loses its geometrical meaning and starts to track the zero-point of the orbit like T_p. Therefore for low eccentricity the relationship between $\omega_{\rm input}$ and $\omega_{\rm output}$ is lost.

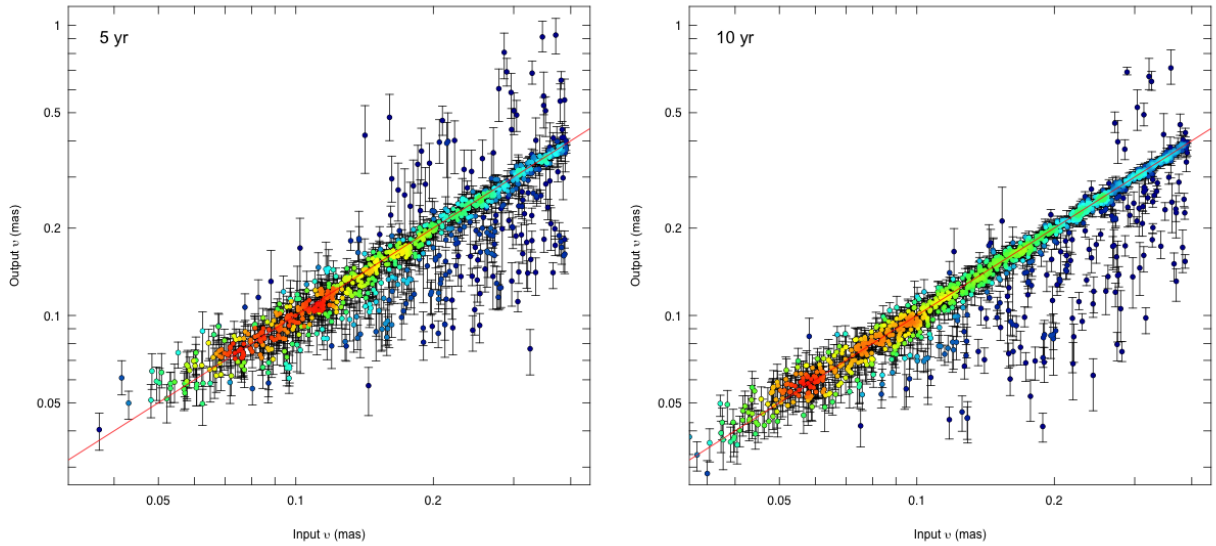

Fig. 10. Simulated vs. recovered angular size of the orbit. Panels, symbols and colours as in Fig. 9. Most of the points with $\upsilon_{\rm output} \ll \upsilon_{\rm input}$ have input periods longer than the mission length, and form the horizontal plume at large $P_{\rm input}$ in Fig. 9.

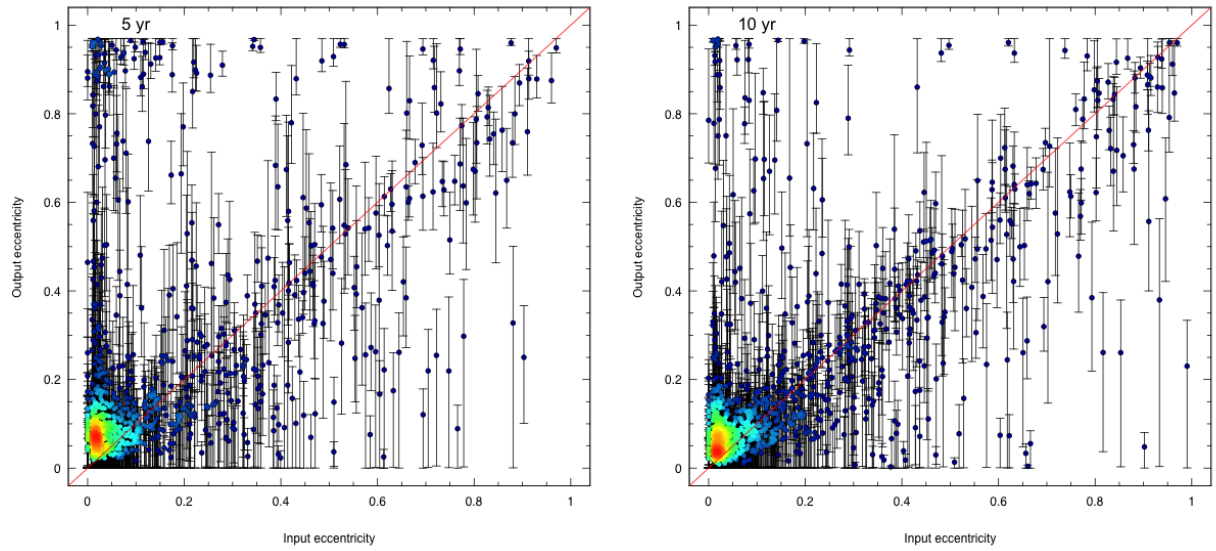

Fig. 11. Simulated vs. recovered eccentricity. Panels, symbols and colours as in Fig. 9. The eccentricity is recovered, albeit with large uncertainties, for most systems with intermediate input eccentricity ($0.1 \lesssim e_{\rm input} \lesssim 0.8$). A number of systems is incorrectly fit with very elliptical orbits ($e_{\rm output} \sim 1$), especially in the 5 yr mission, and corresponds to the parallel sequences of Fig. 9.

The relationship is recovered for larger eccentricities; we plot the points with $\omega_{\rm input} > 0.1$ in Fig. 14. The relationship for ω appears somewhat noisier than the analogous relationships that hold for Ω and T_p.

As discussed above, a small number of systems with incorrect $e_{\rm output} \sim 1$ also have $\omega_{\rm output} = \pm 90°$; however, the error bars in Fig. 14 make these points difficult to discern. To better show them, we plot the subset of points with $0,1 < e_{\rm input} < 0.8$ and $e_{\rm output} > 0.8$ for the 5 yr mission in Fig. 15.

Periapsis transit time

As already seen in the case of ω, the periapsis transit time T_p is only recovered when orbits are not circular. Therefore in Fig. 16 we plot the periapsis transit time only for the systems with $e_{\rm input} > 0.1$. The one-to-one relationship seems well recovered, with no particular structure in the points deviating from it.

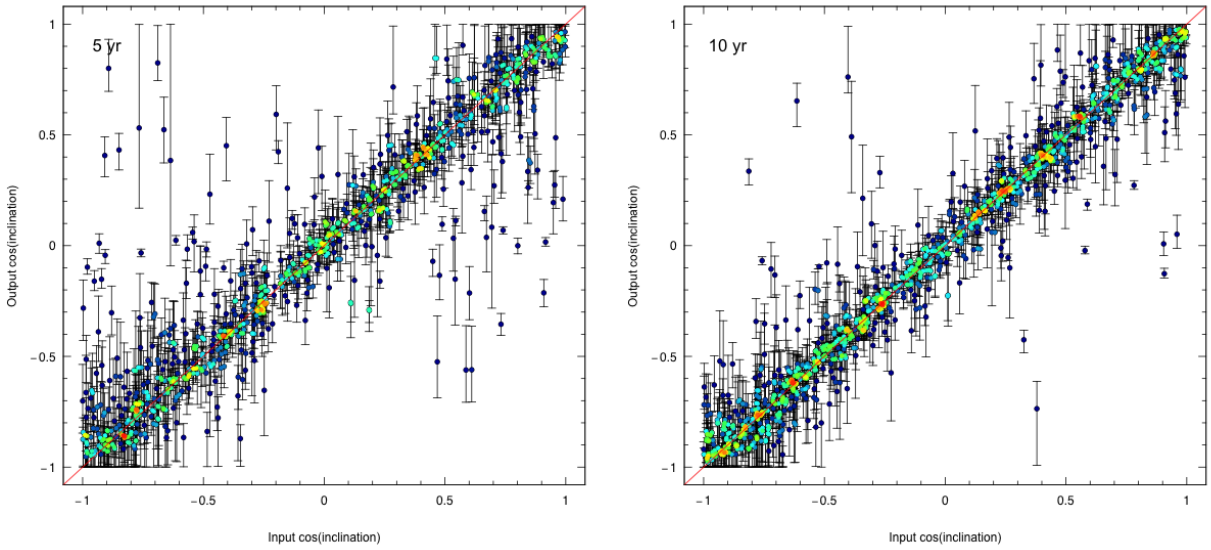

Fig. 12. Simulated vs. recovered inclination. Panels, symbols and colours as in Fig. 9. The inclination is correctly recovered for most systems. Especially in the 5 yr mission, there is a fraction of systems which are mistaken as edge-on ($\cos i_{\rm output} \sim 0$) which seems to be independent of the input inclination but which is often associated with high eccentricity ($e_{\rm output} \sim 1$) and with the parallel sequences in the period.

7. Discussion

7.1. Maximum distance of detectable planets

Previous studies (C08, P14) had suggested a minimum SN = 3 for the detection of planets. This corresponds, according to Fig. 7, to a 50% chance of detecting a planet in the 5 yr mission. Some planets will however be detected even at lower SN. Fig. 7 shows that, for the 5 yr mission and using the WAIC, planets with SN = 2 can be detected in 20% of cases, and planets with SN = 1.5 in 8% of cases. For the 10 yr mission, the numbers rise to 43% and 22%, respectively.

We can therefore take a second look at Fig. 2, where we had presented the SN of a few representative classes of planets, as a function of their distance from the Solar system. Allowing a minimum SN in the 1.5–2 range, the maximum distance up to which a Jovian planet around a 1 $M_\odot$ star could be detected in the 5 yr mission is $\sim$ 50–60 pc, while a Neptunian planet could be detected up to $\sim$ 2–3 pc.

However, in the case of an extended 10 yr mission (that allows longer and wider orbits, up $\sim$ 4 au), and considering lower mass stars (say, 0.5 $M_\odot$), the maximum distance could increase to 120–160 pc for a Jovian planet, and 6–8 pc for a Neptunian.

7.2. Size of errors on fit parameters

The plots in Figs. 9–14 and 16 exhibit a wide range of error bar sizes. It is interesting to examine what parameters are recovered better, and to which accuracy they are on average recovered. We consider the absolute error, defined for each parameter and for each system as the length of the 68.3% HPDI (i.e. the length of error bars in Figs. 9–14 and 16). For two quantities that vary over several orders of magnitude (P and υ) we consider the relative error, defined as the absolute error over the median value of the parameter.

Table 5. Distribution of the errors on the fit parameters. The first column presents the orbital parameter; the other columns report the distribution of the length of its 68.3% error bars in the 5 yr mission (second to fourth column) and in the 10 yr mission (fifth to seventh column) as the 25% quartile, median, and 75% quartile. Relative errors are shown for quantities that vary over several orders of magnitude: P and υ; absolute errors for the remaining parameters. For example, the period (P) has a relative error smaller than ±0.41% in 25% of the cases in the 5 yr mission. We consider T_p/P (instead of just T_p) because it is always contained in the [0, 1] interval; thus for example T_p has an error lower than $\pm 4.8\% P$ in 25% of the cases for the 5 yr mission.

	5 yr			10 yr		
	25%	50%	75%	25%	50%	75%
P	0.41%	1.3%	4.8%	0.18%	0.77%	3.5%
υ	11%	17%	25%	7.2%	13%	20%
e	0.094	0.14	0.21	0.075	0.12	0.18
$\cos i$	0.10	0.16	0.25	0.076	0.13	0.21
ω	7.6°	13°	28°	6.0°	11°	24°
Ω	23°	49°	98°	21°	42°	86°
T_p/P	0.048	0.10	0.23	0.044	0.088	0.19

In this way we obtain for each parameter a distribution of the relative error that it takes for any simulated planetary system; the medians of such distributions may be taken as the characteristic error for each parameter. In Table 5 we show the characteristic (median) relative errors, along with the 25% and 75% quartiles. For example, the period (P) has a relative error smaller than ±0.41% in 25% of the cases in the 5 yr mission.

A more detailed with of the distribution of relative errors is presented in Fig. 17, where we show the error distributions in the form of a violin plot. In this kind of plot, a kernel density estimate is calculated for each parameter distribution, and plotted twice, to form a symmetric figure

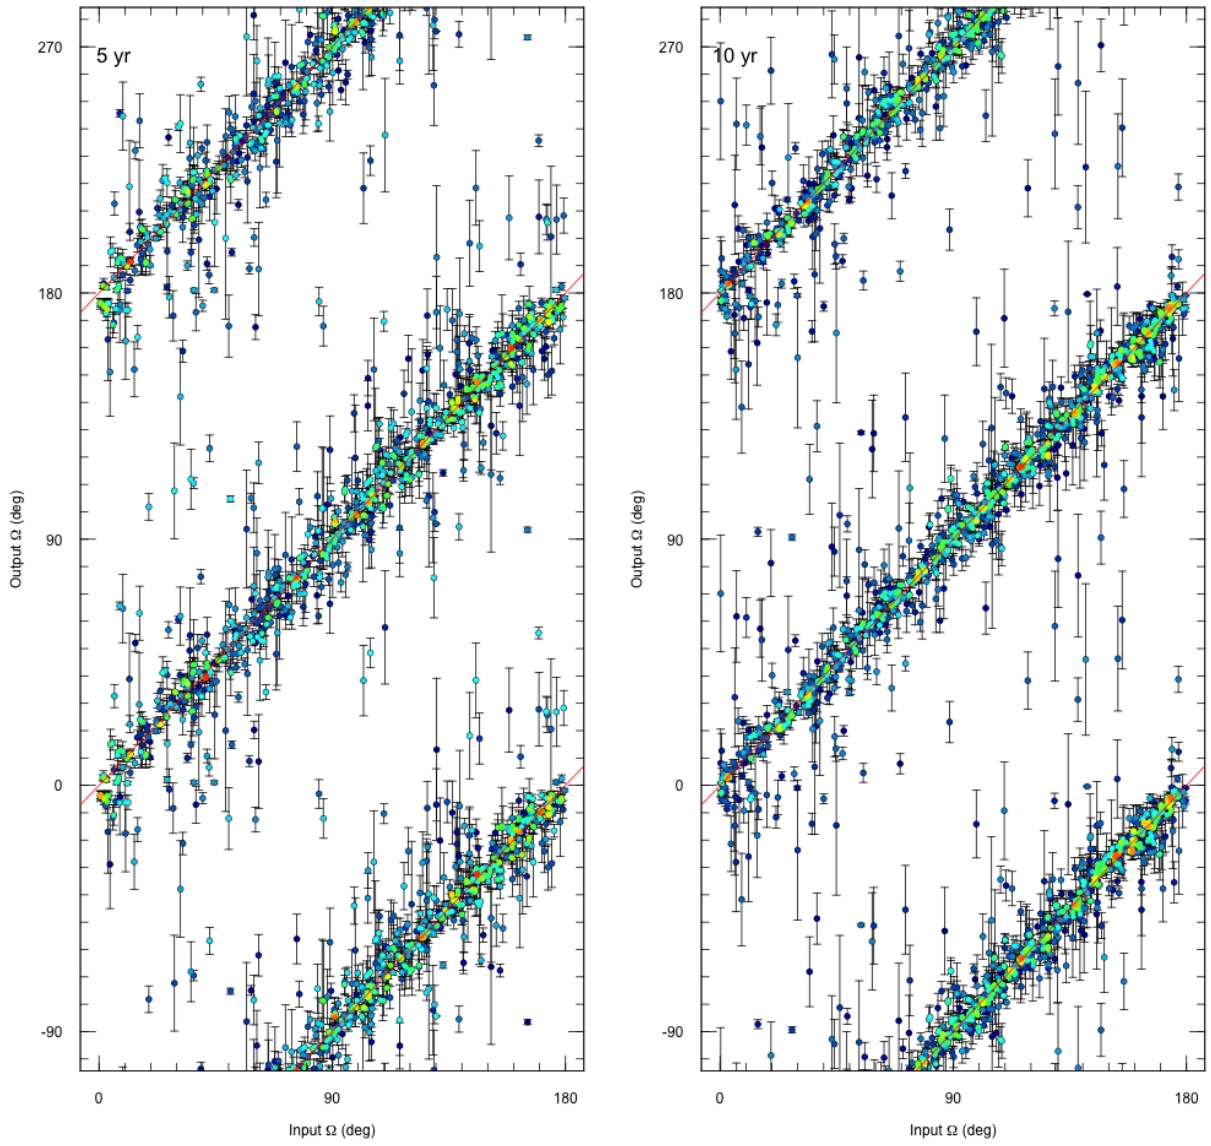

Fig. 13. Simulated vs. recovered longitude of ascending node (Ω). Panels, symbols and colours as in Fig. 9. Since Ω is an angle, we plot the data points twice (at nominal position, and at ±180°) to better show the structure of the noise.

with normalised area. At every value of the abscissa, the figure height is proportional to the density of the distribution.

7.3. Non-converging chains

So far we have considered only systems for which the MC converged, as described in Sect. 5.4. Here we look at the systems for which the MC did not converge. On average, the fraction of converging chains is 80% for the 5 yr mission, and 72% for the 10 yr mission, with no clear dependence on the SN or on any of the fit parameters.

It is not clear how convergence could be improved. One could of course use more burn-in iterations; and in fact some tests showed that using ∼50,000 iterations would improve the convergence fraction, though never reaching fractions close to 100%. More burn-in iterations have however the disadvantage of much longer running times. Since we found no dependence of the converging fraction on any parameter, we deemed the scheme described in Sect. 5.4 as adequate for our needs, and we avoided resorting to more burn-in iterations. For the same reason, we avoided looking deeply into other ways to improve convergence.

However, to apply our methodology to real data, it would be useful to improve the convergence fraction. It is possible that using a different, possibly more complex scheme for the starting point of the MC from that described in Sect. 5.3, might lead to improved convergence. Another possibility could be fine-tuning the MCMC sampler (`Stan` uses the no U-turn sampler, Hoffman & Gelman 2014) or switch to a different sampler such as Gibbs or Metropolis-Hastings.

It is also interesting to consider what information can be extracted from non-converging chains, and if a detection can be assessed despite non-convergence. We found that the WAIC performs differently to the AIC and BIC when chains

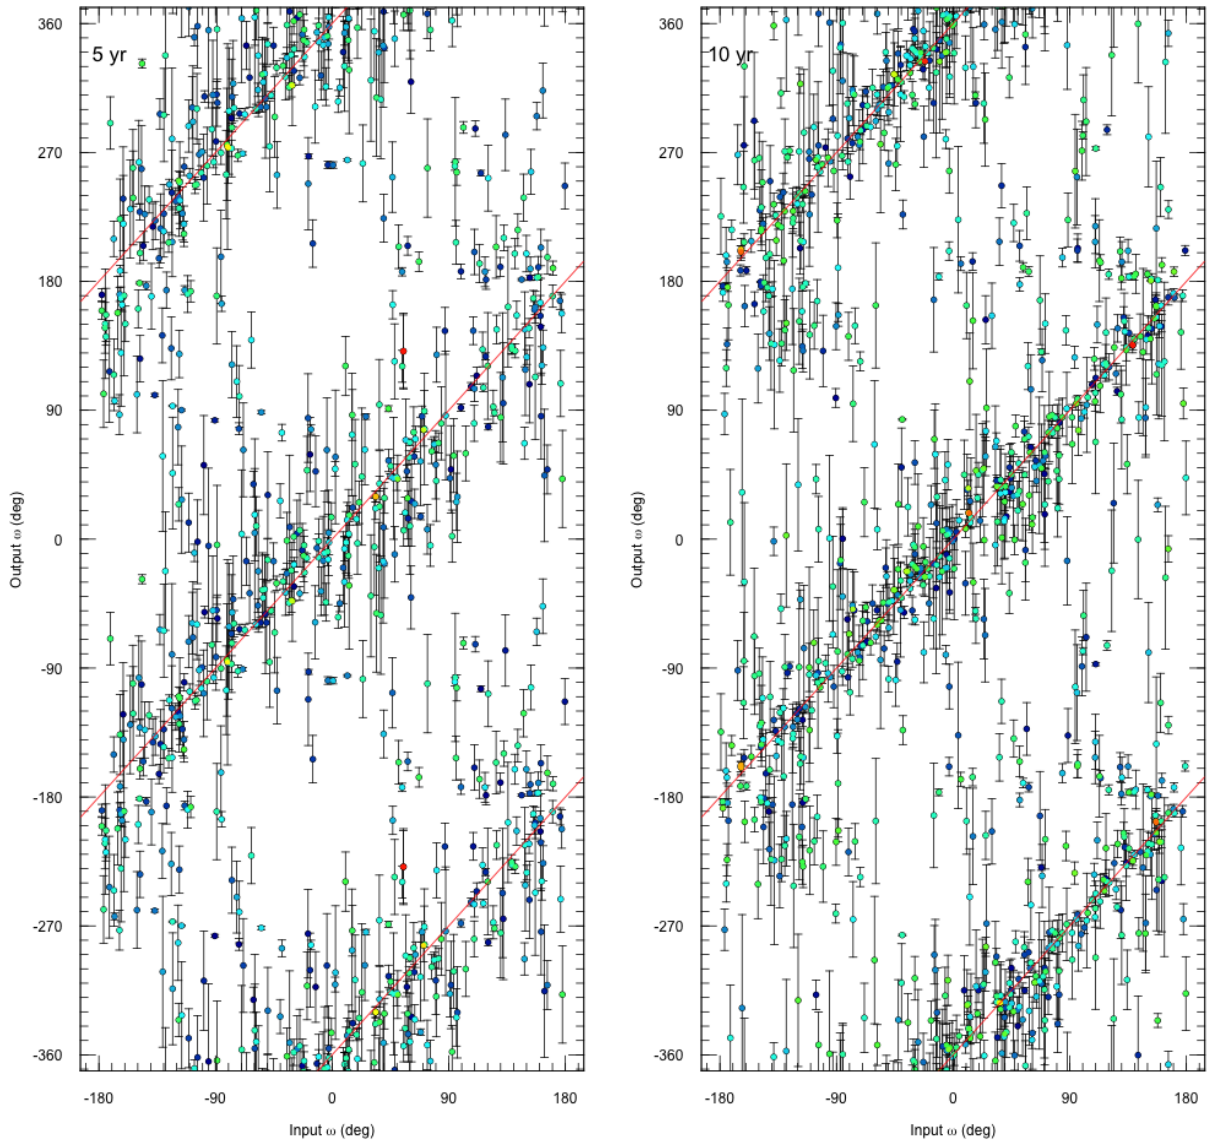

Fig. 14. Simulated vs. recovered argument of periapsis (ω) for systems with $e > 0.1$. Panels, symbols and colours as in Fig. 9. Since ω is an angle, we plot the data points twice (at nominal position, and at $\pm 360°$) to better show the structure of the noise.

do not converge. For systems with SN > 3, the WAIC, AIC, and BIC would signal a detection in 21%, 86%, and 73% of the cases, respectively. Therefore, the WAIC seems to be much more sensitive to convergence problems, thanks to its dependence on the likelihood variance.

We investigated how the parameters were recovered in two groups of systems with non-converging chains: the first group included systems for which the WAIC would signal a detection; the second group included systems for which the AIC or the BIC would signal a detection but the WAIC would not. We found that in the first group the parameters were well recovered; while the second group was dominated by systems which belong to the sort of systems which display an incorrect edge-on, highly-eccentric geometry that we identified in Sect. 6.4.

In Fig. 18 we plot the period, eccentricity and inclination for systems in both groups. The first groups follows the one-to-one relationships much better than the second group: it seems that the WAIC can identify good fits among the non-converging systems. Therefore, future applications may consider the WAIC to be more suited than the AIC or BIC if the fraction of non-converging systems is not negligible.

7.4. Relative efficiency of the information criteria and $\Delta\chi^2$

Our comparison of the different IC in Sects. 6.1 and 6.2 was made on the premise of using a single threshold for all IC. This makes the IC have different detection efficiency for system with intermediate SN (Fig. 5). However, by tuning the threshold for each IC and also for the $\Delta\chi^2$, all detection metrics can achieve a similar efficiency.

In Fig. 19 we show the fraction of detected systems in the 5 yr mission, assuming thresholds equal to the 99.73% quantiles from Table 4. All four detection metrics have a very similar efficiency. This can be explained because all of

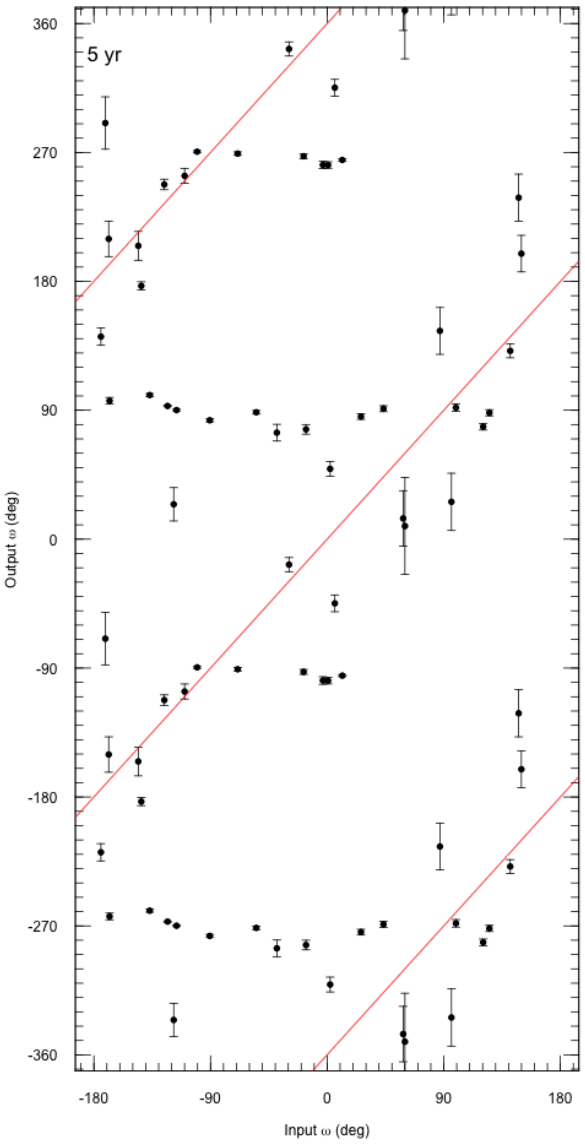

Fig. 15. Simulated vs. recovered argument of periapsis (ω) for systems with $0.1 < e_{\rm input} < 0.8$ but which are incorrectly fit with $e_{\rm output} > 0.8$, for the 5 yr mission. Symbols as in Fig. 9. Since ω is an angle, we plot the data points twice (at nominal position, and at $\pm 360^\circ$) to better show the structure of the noise.

them rely on the likelihood as their basis: the maximum value of the likelihood, for the AIC, BIC, and $\Delta\chi^2$; and the average likelihood for the WAIC. This also explains why the WAIC exhibits somewhat more variability than the other metrics, especially for $4 \lesssim {\rm SN} \lesssim 6$. If detected systems were selected according to these thresholds, a fixed fraction of 0.07% of false positive detections should be expected.

The BIC has a negative 99.73% quantile, that formally should be interpreted as evidence in favour of the null-model (absence of planet). Therefore we have considered a second threshold for the BIC, that we set at $\Delta{\rm BIC}= 0$, and that we show as the brown dot-dashed curve marked as "BIC positive" in Fig. 19. This has the effect of lowering a bit the BIC efficiency.

8. Conclusions

We have presented an investigation of the planet detection capabilities of Gaia, using the most up to date in-flight properties of the spacecraft. We have developed novel methods for the analysis of astrometric data of planetary orbits, and for the detection of planets. Our methods operate in the Bayesian framework, relying on well-established software and procedures for inference and model selection.

We have presented two models, one accounting only for astrometric parameters (position, parallax, proper motion), and another accounting for both astrometric and orbital parameters; a scheme for detecting convergence of the Markov chains; and three information criteria to be used for model selection (AIC, BIC, WAIC).

Using two set of simulations of Gaia observations of stars with no planets (one set for a nominal 5 yr mission, and another set for an extended 10 yr mission), we have compared the information criteria against the threshold of $\Delta{\rm IC} = 20$ suggested for strong evidence in favour of one model over another by Jeffreys' scale for Bayesian evidence. We have found that the above threshold allows to avoid false detections in more than 99.7% of cases, with the AIC being more conservative than the WAIC, and BIC more conservative then the AIC and WAIC.

Using the same sets of simulations, we have showed that the $\Delta\chi^2$ can also be used as a detection indicator, provided that a $\Delta\chi^2$ threshold is empirically calibrated, because the non-linearity of the problem of orbit modelling makes the effective number of degrees of freedom a data-dependent variable, whose value we found to depend on the mission length (because of the larger number of data points collected in an extended mission).

Our main sets of simulations include stars with one planet, whose SN, periods, and eccentricities were assigned according to a logarithmically-spaced grid. All other parameters were assigned randomly. Realistic scanning laws were calculated according to a randomly assigned position on the sky.

Applying our detection method to the simulations of stars with one planet, we have found that a threshold of $\Delta{\rm IC} = 20$, while not including false detections, was allowing a number of low-quality detections with fit parameter not matching the input ones. Therefore we have adopted a more conservative threshold of $\Delta{\rm IC} = 70$, which improved the quality of fits.

We have confirmed the finding of previous studies that the SN is the main parameter that determines whether a planet can be detected. While we confirm the rule-of-thumb threshold of ${\rm SN} \gtrsim 3$, we have also found that a ${\rm SN} = 1.5$ allows detection in 8% and 22% of cases for a 5 yr and 10 yr mission, respectively. This means that the approximate maximum distance to which a planet may be detected is 60 pc (160 pc) for a Jovian planet and 3 pc (8 pc) for a Neptunian planet, in the case of a 5 yr (10 yr) mission.

The orbital parameters of planets with periods up to the mission length can be recovered with different degrees of accuracy. For exampl, for a 5 yr mission we find a median relative error of $\pm 1.3\%$ for the period; or $\pm 17\%$ for the angular size of the semi-major axis of the orbit; or a median absolute error of ± 0.14 for the eccentricity. Also, in the case of nearly-circular orbits, a degeneracy between the argument of periapsis and the periapsis transit time hinders the identification of both; in those cases, a different model

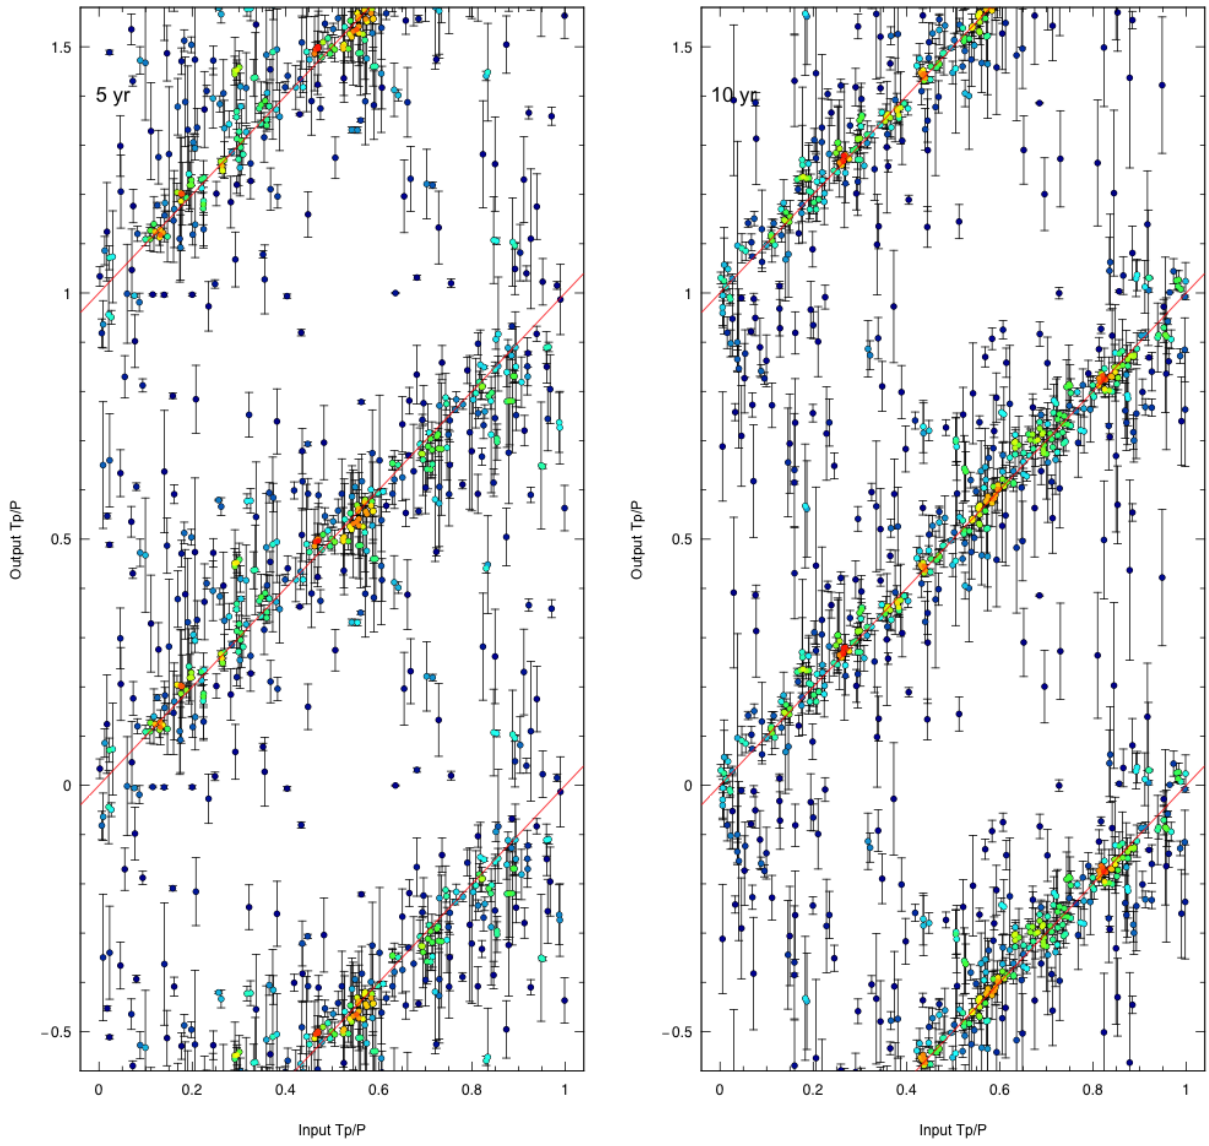

Fig. 16. Simulated vs. recovered periapsis transit time (Tp) normalised by the period, for systems with $e > 0.1$. Panels, symbols and colours as in Fig. 9. Since Tp/P has a circular support, we plot the data points twice (at nominal position, and at ± 1) to better show the structure of the noise.

which fixes one of those parameters, or explicitly considers circular orbits, may be preferred.

Considering planets with periods longer than the mission length, we have found that they can be detected, provided that the SN is high enough (SN $\gtrsim 6$ for periods twice the mission length), albeit the recovered period will not match the real one, but will rather stand at $\sim 80\%$ of the mission length.

Finally, we have considered the systems for which the Markov chains did not converge: we have found that while the three information criteria do in some cases signal a detection, only the WAIC is able to identify the systems for which the orbital parameters can be reliably obtained.

A follow-up of this work should concern the detectability of multiple planets in the same systems, and of very long-period planets from using data from both Gaia and a future astrometric mission.

References

Akaike, H. 1973, in 2nd International Symposium on Information Theory, Academiai Kiado, reprinted in Selected Papers of Hirotugu Akaike (Springer 1998), 199–213
Andrae, R., Schulze-Hartung, T., & Melchior, P. 2010, ArXiv e-prints
Binnendijk, L. 1960, Properties of double stars : a survey of parallaxes and orbits. (Philadelphia : Pennsylvania U.P., 1960)
Burnham, K. P. & Anderson, D. R. b. 2002, Model selection and multimodel inference : a practical information-theoretic approach (New York: Springer)
Carpenter, B., Gelman, A., Hoffman, M., et al. 2016, Journal of Statistical Software, 20
Casertano, S., Lattanzi, M. G., Sozzetti, A., et al. 2008, A&A, 482, 699
de Bruijne, J., Siddiqui, H., Lammers, U., et al. 2010, in IAU Symposium, Vol. 261, Relativity in Fundamental Astronomy: Dynamics, Reference Frames, and Data Analysis, ed. S. A. Klioner, P. K. Seidelmann, & M. H. Soffel, 331–333
de Bruijne, J. H. J. 2012, Ap&SS, 341, 31
de Bruijne, J. H. J., Rygl, K. L. J., & Antoja, T. 2014, in EAS Publications Series, Vol. 67, EAS Publications Series, 23–29

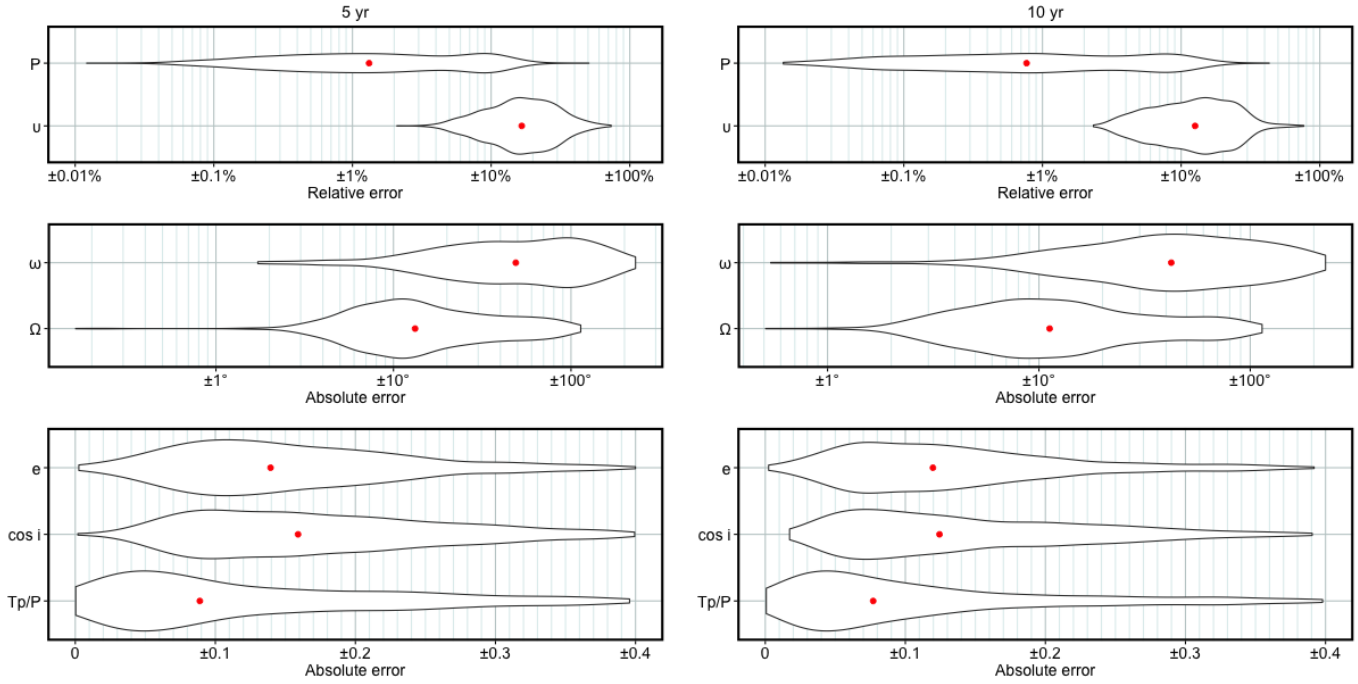

Fig. 17. Size of errors on the fit parameters. Left panels: 5 yr mission; right panels: 10 yr mission. For each orbital parameter, we plot the distribution of the length of the 68.3% highest posterior density interval (HPDI) . The red dots mark the median of each distribution. The figures depicted by the black curves have normalised areas; at every value of the abscissa, their height is proportional to the density of the distribution. Relative errors are shown for quantities that vary over several orders of magnitude: P and υ; absolute errors for the remaining parameters. We consider T_p/P (instead of just T_p) because it is always contained in the [0, 1] interval; thus for example T_p has a median error of $5\%P$ in the 5 yr mission.

Efron, B. & Gous, A. 2001, in Lecture Notes-Monograph Series, ed. P. Lahiri, Vol. 38: Model Selection (Institute of Mathematical Statistics), pp. 208–256

ESA, ed. 1997, ESA Special Publication, Vol. 1200, The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission

Gaia Collaboration. 2016, ArXiv e-prints

Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2016, A&A, 595, A2

Gatewood, G., Han, I., & Black, D. C. 2001, ApJ, 548, L61

Gelman, A., Carlin, J., Stern, H., et al. 2013, Bayesian Data Analysis, Third Edition, Chapman & Hall/CRC Texts in Statistical Science (Taylor & Francis)

Gelman, A. & Rubin, D. B. 1992, Statistical Science, 7, 457

Geyer, C. 2011, in Handbook of Markov Chain Monte Carlo, ed. S. Brooks, A. Gelman, G. L. Jones, & X.-L. Meng, Chapman & Hall/CRC Handbooks of Modern Statistical Methods (Chapman and Hall/CRC), 3–48

Gregory, P. 2005, Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica® Support (Cambridge University Press)

Han, I., Black, D. C., & Gatewood, G. 2001, ApJ, 548, L57

Hoffman, M. D. & Gelman, A. 2014, Journal of Machine Learning Research, 15, 1593

Holl, B., Lindegren, L., & Hobbs, D. 2012, A&A, 543, A15

Jacob, W. S. 1855, MNRAS, 15, 228

Jammalamadaka, S. & SenGupta, A. 2001, Topics in Circular Statistics, Series on multivariate analysis (World Scientific)

Liddle, A. R. 2007, MNRAS, 377, L74

Lindegren, L., Lammers, U., Bastian, U., et al. 2016, A&A, 595, A4

Lindegren, L., Lammers, U., Hobbs, D., et al. 2012, A&A, 538, A78

Lunn, D. J., Thomas, A., Best, N., & Spiegelhalter, D. 2000, Statistics and computing, 10, 325

Markley, F. L. 1995, Celestial Mechanics and Dynamical Astronomy, 63, 101

Mazeh, T., Zucker, S., Dalla Torre, A., & van Leeuwen, F. 1999, ApJ, 522, L149

McArthur, B. E., Benedict, G. F., Barnes, R., et al. 2010, ApJ, 715, 1203

Michalik, D., Lindegren, L., Hobbs, D., & Lammers, U. 2014, A&A, 571, A85

Mikkola, S. 1987, Celestial Mechanics, 40, 329

Perryman, M., Hartman, J., Bakos, G. Á., & Lindegren, L. 2014, ApJ, 797, 14

Perryman, M. A. C. 2011, The exoplanet handbook. (Cambridge ; New York : Cambridge University Press)

Pewsey, A., Neuhäuser, M., & Ruxton, G. 2013, Circular Statistics in R, EBL-Schweitzer (OUP Oxford)

Plummer, M. et al. 2003, in Proceedings of the 3rd international workshop on distributed statistical computing, Vol. 124, Vienna, 125

Pourbaix, D. 2001, A&A, 369, L22

Quirrenbach, A. 2010, in Exoplanets, ed. S. Seager (Tucson ; The University of Arizona press), 157–174

Ranalli, P., Koulouridis, E., Georgantopoulos, I., et al. 2016, A&A, 590, A80

Schwarz, G. 1978, Ann. Statist., 6, 461

Trotta, R. 2008, Contemporary Physics, 49, 71

Watanabe, S. 2010, Journal of Machine Learning Research, 11, 3571

Watanabe, S. 2013, Journal of Machine Learning Research, 14, 867

Zucker, S. & Mazeh, T. 2000, ApJ, 531, L67

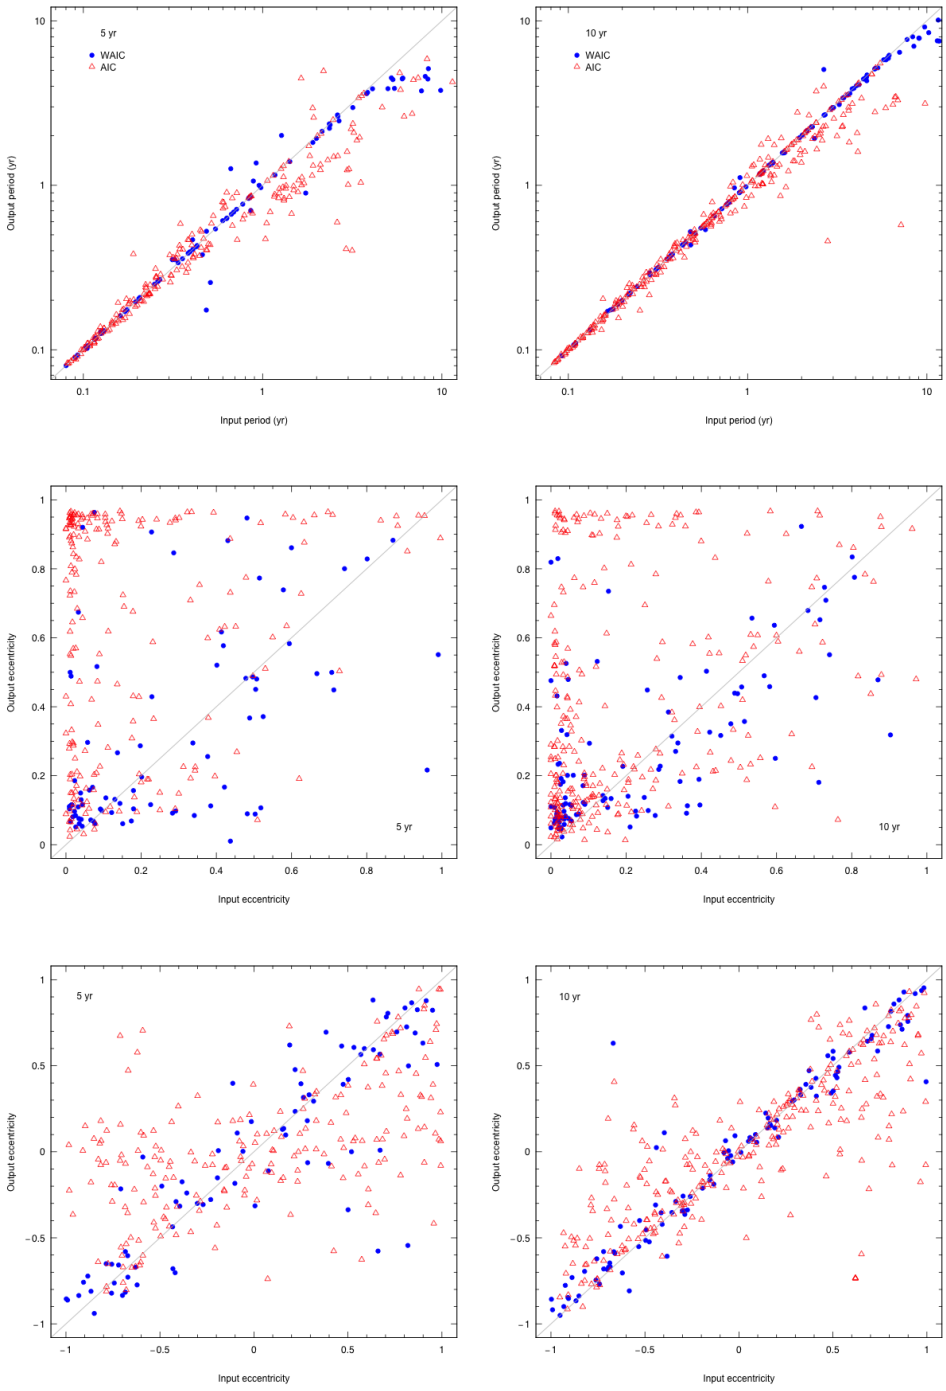

Fig. 18. Period, eccentricity, and inclination for systems whose MC did not converge. Blue filled circles: systems detected according to the WAIC; red open triangles: systems detected according to the AIC but not according to the WAIC. The WAIC can identify good fits among the non-converging systems better than the AIC.

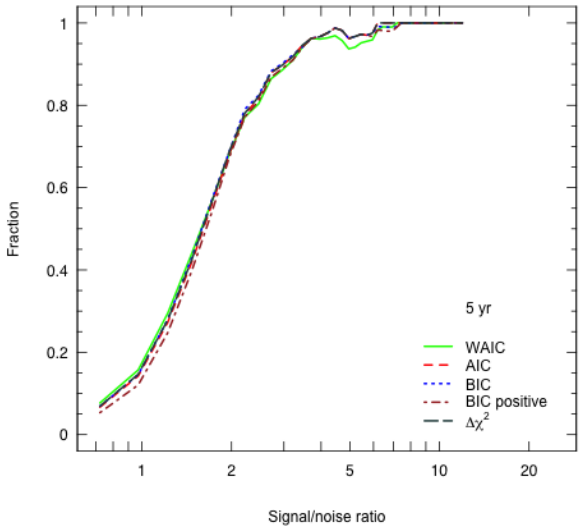

Fig. 19. Fraction of detected systems as a function of SN, with detection thresholds differing for each criterion, and taken equal to the 99.73% quantiles from Table 4. In this way, all criteria achieve the same detection efficiency. However, the ΔBIC threshold gets set to a negative value, which formally means that the planet cannot be considered detected. Therefore we have also added the brown dot-dashed line to show the case of ΔBIC> 0.